



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ :

G06T 9/00, 9/40

A1

(11) International Publication Number:

WO 97/32281

(43) International Publication Date: 4 September 1997 (04.09.97)

(21) International Application Number: PCT/US97/03300

(22) International Filing Date: 27 February 1997 (27.02.97)

(30) Priority Data:

08/607,388

27 February 1996 (27.02.96)

US

(60) Parent Application or Grant

(63) Related by Continuation

US

Filed on

08/607,388 (CIP)

27 February 1996 (27.02.96)

(71) Applicant (for all designated States except US): INTERVAL RESEARCH CORPORATION [US/US]; Building C, 1801 Page Mill Road, Palo Alto, CA 94304 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): KOLAROV, Krasimir, D. [BG/US]; 2050 Avy Avenue, Menlo Park, CA 94025 (US). LYNCH, William, C. [US/US]; 3331 Thomas Drive, Palo Alto, CA 94303 (US). SCHRÖDER, Peter [DE/US]; 1111 Blanche Street #201, Pasadena, CA 91106 (US). SWELDENS, Wim [BE/US]; Apartment 2B, 23 Euclid Avenue, Summit, NJ 07901 (US).

(74) Agents: WEAVER, Jeffrey, K. et al.; Hickman Beyer & Weaver, P.O. Box 61059, Palo Alto, CA 94306 (US).

(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, HU, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PT, RO, RU, SD, SE, SG, SI, SK, TJ, TM, TR, TT, UA, UG, US, UZ, VN, ARIPO patent (GH, KE, LS, MW, SD, SZ, UG), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, SJ, TM), European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).

Published


With international search report.

Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.

(54) Title: WAVELET BASED DATA COMPRESSION

(57) Abstract

A technique for compression and expansion of a function defined upon an M-dimensional manifold embedded in N-dimensional space uses a second generation wavelet transform and a modified zerotree bit-encoding scheme. Typically, a function is defined upon a two-dimensional manifold embedded in three-dimensional space, such as a sphere. A geometric base is chosen as a coarse initial model of the manifold. Second generation wavelets for the function are calculated using a triangular subdivision scheme in order to subdivide the geometric base in order to produce a refined triangular mesh. The wavelet coefficients are defined at the vertices of the triangles in the triangular mesh. A tree structure is created in which each node of the tree structure represents an associated triangle of the triangular mesh. Each triangle in the mesh is recursively subdivided into four subtriangles and each associated node in the tree structure also has four children, which correspond to the four subtriangles. Each wavelet coefficient defined at a particular vertex in triangular mesh is uniquely assigned to a single one of the triangles at a next higher level of subdivision, such that each triangle at the next higher level of subdivision has from zero to three assigned wavelet coefficients. Using a modified zerotree encoding scheme, values of the wavelet coefficients are processed bit plane by bit plane, outputting bits indicative of significant nodes and their descendants. Sign bits and data bits are also output. An expansion technique inputs bits according to the modified zerotree scheme into the tree structure in order to define wavelet coefficients. An inverse second generation wavelet transform is used to synthesize the original function from the wavelet coefficients.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AM	Armenia	GB	United Kingdom	MW	Malawi
AT	Austria	GE	Georgia	MX	Mexico
AU	Australia	GN	Guinea	NE	Niger
BB	Barbados	GR	Greece	NL	Netherlands
BE	Belgium	HU	Hungary	NO	Norway
BF	Burkina Faso	IE	Ireland	NZ	New Zealand
BG	Bulgaria	IT	Italy	PL	Poland
BJ	Benin	JP	Japan	PT	Portugal
BR	Brazil	KE	Kenya	RO	Romania
BY	Belarus	KG	Kyrgyzstan	RU	Russian Federation
CA	Canada	KP	Democratic People's Republic of Korea	SD	Sudan
CF	Central African Republic	KR	Republic of Korea	SE	Sweden
CG	Congo	KZ	Kazakhstan	SG	Singapore
CH	Switzerland	LI	Liechtenstein	SI	Slovenia
CI	Côte d'Ivoire	LK	Sri Lanka	SK	Slovakia
CM	Cameroon	LR	Liberia	SN	Senegal
CN	China	LT	Lithuania	SZ	Swaziland
CS	Czechoslovakia	LU	Luxembourg	TD	Chad
CZ	Czech Republic	LV	Latvia	TG	Togo
DE	Germany	MC	Monaco	TJ	Tajikistan
DK	Denmark	MD	Republic of Moldova	TT	Trinidad and Tobago
EE	Estonia	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	UG	Uganda
FI	Finland	MN	Mongolia	US	United States of America
FR	France	MR	Mauritania	UZ	Uzbekistan
GA	Gabon			VN	Viet Nam

WAVELET BASED DATA COMPRESSION

FIELD OF THE INVENTION

5 The present invention relates generally to the compression of data defined on complex geometries. More specifically, a compression technique that is adapted to compress data defined using a wavelet transformation is described.

BACKGROUND OF THE INVENTION

10 Data compression techniques are useful to represent information as accurately as possible with a minimum number of bits in order to minimize the amount of data that must be stored or transmitted for use in a computer system. Typically, data to be compressed may represent various types of information, such as audio and video information. One of the means primarily used for compressing data is the removal of redundant information. At the same time, however, the remaining information must accurately represent the data that has been compressed.

15 Traditional data compression methods work on data defined on simple geometries. Typical geometries used in conventional compression methods include: the real line (e.g., audio), a rectangle (e.g., images), or a three-dimensional raster (e.g., video). However, these conventional data compression technologies are not suitable for compression of data defined on higher-dimensional or complex geometries such as spheres, general polygonal models, etc.

20 There are numerous examples of complex geometries for which it would be useful to develop an appropriate data compression technique. One example is Earth data compression. At the moment very large data sets of Earth elevation data exists. A typical example is the ETOPO5 data set. This data set represents topographic data over the entire Earth. The data gives the elevation (depth) of the Earth from sea level in meters at a resolution of 5 arc minutes at the equator. Even with powerful workstations, manipulating, storing, and transmitting such data sets
25 can be quite expensive. Other examples of similar data sets are satellite views of the face of the Earth and planetary data. Another example of a complex geometry that is useful for modeling information is the human head. Another relevant application is the interactive multi-resolution viewing of complex data sets. When building an interactive viewer for complex manifolds, it would be desirable to only store and display the information needed by the user.

30 While various techniques are available in the prior art for compression of the simple geometries, no unified technique currently exists for the compression of data defined on complex and/or higher dimensional geometries. Accordingly, the present invention seeks to provide such a technique.

One definition of a wavelet is: "Wavelets are building blocks that can quickly decorrelate data." The wavelet representation of a data set can be found quickly. More precisely, one can switch between the original representation of the data and its wavelet representation in a time proportional to the size of the data. The fast decorrelation power of wavelets is one key to compression applications.

In general, wavelets are basis functions that represent a given function at multiple levels of detail. Due to their local support in both space and frequency, they are suited for the sparse approximations of functions. Locality in space follows from their compact support, while locality in frequency follows from their smoothness and vanishing moments. In the classic wavelet setting, i.e., on the real line, wavelets are defined as the dyadic translates and dilates of one particular, fixed function. They are typically built with the aid of a *scaling function*. Scaling functions and wavelets both satisfy refinement relations (or two scale relations). This means that a scaling function or wavelet at a certain level of resolution (j) can be written as a linear combination of scaling basis functions of the same shape but scaled at one level finer (level $j + 1$).

Over the last decade wavelets have become an exceedingly powerful and flexible tool for computations and data reduction. They offer theoretical characterizations of smoothness, insights into the structure of functions and operators, and practical numerical tools that lead to faster computational algorithms. Examples of their use in computer graphics include surface and volume illumination computations, curve and surface modeling and animation, among others. Given the high computational demand in the quest for speed in computer graphics, the increasing exploitation of wavelets comes as no surprise. While these and other applications can benefit greatly from wavelet techniques, some of these applications can provide new challenges to the underlying wavelet technology. One such challenge is the transformation of data defined on complex manifolds. In general, an M -dimensional manifold is a space N such that each point in N has a neighborhood homeomorphic to an open ball in R to the M power. For example, a two-dimensional (2-D) spherical surface embedded in Euclidean three-dimensional (3-D) space is such that each small portion of the sphere appears to be the same as a small portion of Euclidean 2-D space (i.e., the plane). This fact caused some people to believe that the Earth was flat.

Classic wavelet constructions define the wavelet functions as dyadic translates and dilates of one fixed function: "the mother wavelet." See, for example, C.K. Chui, "An Introduction to Wavelets," Academic Press, 1992; A. Cohen, I. Daubechies, and J. Feauveau, "Bi-orthogonal Bases of Compactly Supported Wavelets," Comm. Pure Appl. Math., 45:485-560, 1992; and I. Daubechies, "Orthonormal bases of compactly supported wavelets," Comm. Pure Appl. Math., 41:909-996, 1988. Because translation and dilation only make sense in a flat Euclidean space, these classical, or first generation wavelets cannot be used for the transformation of data defined on more complex manifolds. Therefore, one needs to consider the generalization to *second generation wavelets*. These are wavelets that are not necessarily the translates and dilates of one

function, but that still enjoy all the powerful properties of first generation wavelets, namely basis properties, decorrelation power, and fast algorithms. In general, it can be said that most all of the widely used classical wavelets may be represented using second generation wavelets, and it is believed that all classical wavelets may be so represented. In other words, second generation wavelets are a superset of most all classical wavelets. Thus, the term "second generation wavelets" as used herein refers not only to the specific second generation wavelets described, but also to those classical wavelets that may be represented using second generation wavelets.

The basic philosophy behind second generation wavelets is to build wavelets with all desirable properties (localization, fast transform) adapted to much more general settings than the real line, e.g., wavelets on manifolds. In order to consider wavelets on a surface, one needs a construction of wavelets that are adapted to a measure on the surface. In the case of the real line (and classical constructions) the measure is dx , the usual translation invariant Lebesgue measure. For a sphere, the usual area measure can be denoted by dw , for example. Adaptive constructions rely on the realization that translation and dilation are not fundamental to obtain the wavelets with the desired properties. The notion that a basis function may be written as a finite linear combination of basis functions at a finer, more subdivided level, is maintained and forms the key behind the fast transform. The main difference with the classical wavelets is that the filter coefficients of second generation wavelets are not the same throughout, but can change locally to reflect the changing (non-translation invariant) nature of the surface and its measure.

Classical wavelets and the corresponding filters are constructed with the aid of the Fourier transform. The underlying reason is that translation and dilation become algebraic operations after the Fourier transform. In the setting of second generation wavelets, translation and dilation can no longer be used, and the Fourier transform cannot be used as a construction tool. An alternative construction is provided by the *lifting scheme*.

This new concept in the construction of wavelets, the lifting scheme, has been introduced recently. This concept is described in W. Sweldens, "The Lifting Scheme: A Custom-design Construction of Biorthogonal Wavelets", Technical Report 1994:7, Industrial Mathematics Initiative, Department of Mathematics, University of South Carolina, 1994, which is incorporated herein by reference in its entirety. This lifting scheme differs from the classical constructions in the sense that it does not need to use Fourier transforms. Therefore, it is an excellent tool for the construction of second generation wavelets. In one implementation, triangular subdivision is used to model a three-dimensional manifold, such as a sphere. Such a triangular subdivision of a sphere 200 is generally illustrated in Figure 2a. In this figure the surface of sphere 200 has been subdivided into triangles in the form of a geodesic sphere. Shown is a triangle 210 that represents one of the many triangles on the surface of the sphere. Note that triangle 210 itself forms a part of a larger triangle that is represented by triangles 211, 210, 212 and 213. It should also be noted that triangle 210 itself may be further subdivided according to the same scheme.

While modeling techniques that utilize second generation wavelets to represent data on complex manifolds have proven useful, working with such data can be extremely difficult since the amount of data required for modeling can be extremely large, and presently, there are no suitable mechanisms for compressing such data. Accordingly, the development of a
5 compression arrangement suitable for compressing such data would be highly desirable.

SUMMARY OF THE INVENTION

The present invention embodies a compression technique for compressing and uncompressing data representing a function defined upon an M-dimensional manifold embedded in N-dimensional space. With an M-dimensional manifold, a geometric base type is used as a
10 coarse initial model of the manifold and the geometric base is subdivided into cells. Second generation wavelet coefficients for the function are generated using a subdivision scheme in order to subdivide the geometric base into cells. The coefficients are defined upon the cells.

The invention may also compress and uncompress a function defined upon a two-dimensional manifold, where the manifold is embedded in three-dimensional space. In one
15 aspect of compression, second generation wavelets for the function are generated using a triangular subdivision scheme in order to subdivide the geometric base in order to produce a refined triangular mesh. The wavelet coefficients are defined at the vertices of the triangles in the triangular mesh. A tree structure is created in which each node of the tree structure represents an associated triangle of the triangular mesh. Each triangle in the mesh is recursively subdivided into
20 four subtriangles and each associated node in the tree structure also has four children, which correspond to the four subtriangles. Each wavelet coefficient defined at a particular vertex in the triangular mesh is uniquely assigned to a single one of the triangles at a next higher level of subdivision, such that each triangle at the next higher level of subdivision has from zero to three assigned wavelet coefficients.

25 In order to compress the wavelet coefficients, the values of the wavelet coefficients are processed bit plane by bit plane, by stepping through the nodes of the tree structure. A first threshold is set, and a stream of bits is output indicative of significant nodes. A significant node is defined as a node that has either an associated wavelet coefficient that exceeds the designated threshold or a descendant node that has an associated wavelet coefficient that exceeds the
30 designated threshold. When a particular significant node is determined, the significant children are recursively examined for significant descendants and the output bit stream identifies the significant descendants and the significant children nodes. After all nodes have been processed in this fashion for this bit plane, a new threshold is set and the next bit plane is processed. The technique continues in this fashion until all bit planes for the wavelet coefficients have been
35 processed.

In a further embodiment, when the sign bit of a wavelet coefficient is encountered, it is also output in the bit stream. Also, in a refinement step, for each node already determined to be significant, and not newly determined to be significant, the bit corresponding to the current bit plane is output for each wavelet coefficient associated with that node. In another embodiment, the output bit stream may be entropy encoded using a technique such as Huffman encoding, binary encoding, arithmetic encoding, run-length encoding, or LZW encoding.

In one embodiment, the manifold is a sphere, and the geometric base type may be an icosahedron, an octahedron, or a tetrahedron. For each of the geometric base types, the calculation of second generation wavelets may use any of a linear, quadratic, cubic or butterfly lifting scheme.

In another embodiment a wavelet coefficient is uniquely associated with a particular triangle according to the following technique. As a coefficient is defined at the vertex of a given triangle, this vertex will be incident upon the midpoint of an edge that is shared by two larger triangles at the next higher level of subdivision. The coefficient will be uniquely associated with one of these two triangles. In one embodiment, all triangles of the mesh are parameterized either clockwise or counter-clockwise, and the coefficient is uniquely associated with one of the larger triangles according to its parametrization.

The present invention also embodies a technique to perform uncompression of a function defined upon a two-dimensional manifold, where the manifold is embedded in three-dimensional space. The invention may also be generalized and apply to an M-dimensional manifold embedded in N-dimensional space. The function may have been compressed using the above described compression technique. The uncompression technique defines a tree structure as described above where each node of the tree structure represents a triangle of the triangular mesh. From zero to three wavelet coefficients may be associated with each triangle in the mesh, although the values for these coefficients will not be assigned until the data is input and uncompressed.

The uncompression technique sets a first threshold and inputs a stream of significance bits indicative of significant nodes of the tree structure. When a particular significant node is determined, the input stream of significance bits further includes bits indicative of the significance of children nodes of the particular significant node and the input bit stream identifies the significant descendants. In addition, also input are a stream of data bits indicative of values for second generation wavelet coefficients, the wavelet coefficients being defined on triangles within the refined mesh. Once the wavelet coefficients have been defined, the technique synthesizes the data representing the function from the wavelet coefficients in order to reconstruct the function defined upon the manifold. A synthesis technique may perform the inverse wavelet transform that was used to calculate the wavelet coefficients.

BRIEF DESCRIPTION OF THE DRAWINGS

The invention, together with further advantages thereof, may best be understood by reference to the following description taken in conjunction with the accompanying drawings in which:

5 Figure 1a is a symbolic representation of an image undergoing a wavelet transform.

Figure 1b is a symbolic representation of the same image in a further iteration of a wavelet transform.

Figure 1c is a symbolic representation of the same image in another iteration of a wavelet transform.

10 Figure 1d shows a tree structure of wavelet coefficients used to implement the zerotree algorithm.

Figure 2a shows a spherical object whose surface has been subdivided into triangles according to a geodesic subdivision scheme.

15 Figure 2b shows a triangle from the surface of the sphere that has undergone a first subdivision as part of a wavelet transform.

Figure 2c shows a triangle from the surface of the sphere that has undergone a second subdivision as part of a wavelet transform.

Figure 3a is a flowchart showing a process by which data representing a function defined upon a complex manifold is compressed according to one embodiment of the present invention.

20 Figure 3b is a flowchart showing a process by which compressed bits in a data file representing a function defined upon a complex manifold are uncompressed according to one embodiment of the present invention.

Figures 4a, 4b and 4c show a flowchart for a process in which wavelet coefficients represented in a tree structure are compressed and output to a file.

25 Figures 5a, 5b and 5c show a flowchart for a process in which compressed data in a file is uncompressed into the original wavelet coefficients.

Figure 6 shows a triangle from the surface of a sphere that is being subdivided in a wavelet transform in accordance with an embodiment of the invention.

30 Figures 7a, 7b and 7c illustrate alternative subdivision schemes for subdividing a geometric shape according to other embodiments of the invention.

Figure 8 shows a computer-based system in accordance with an embodiment of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

The present invention relates to computer readable media that include program instructions for performing various computer-implemented operations. The media and program instructions may be those specially designed and constructed for the purposes of the present invention, or they may be of the kind well known and available to those having skill in the computer software arts. Examples of computer readable media include, but are not limited to, magnetic media such as hard disks, floppy disks, and magnetic tape; optical media such as CD-ROM disks; magneto-optical media such as floptical disks; and hardware devices that are specially configured to store and perform program instructions, such as read-only memory devices (ROM) and random access memory (RAM). Examples of program instructions include both machine code, such as produced by a compiler, and files containing higher level code that can be executed by the computer using an interpreter.

Referring next to Figure 8, a typical computer-based system suitable for implementing the present invention is diagrammatically illustrated. Shown is a central processing unit 702 (CPU) that is coupled to memory devices including read only memory 704 (ROM) and random access memory 706 (RAM). As is well known in the art, ROM 704 acts to transfer data and instructions uni-directionally to the CPU and RAM 706 is used typically to transfer data and instructions in a bi-directional manner. A mass memory device 708 is also coupled bi-directionally to CPU 702 and provides additional data storage capacity. The mass memory device 708 may be used to store programs, data and the like and may take the form of a magnetic or paper tape reader or some other well-known device (*e.g.*, CD-ROM). It will be appreciated that the information retained within the mass memory device 708, may, in appropriate cases, be incorporated in standard fashion as part of RAM 106 as virtual memory. CPU 702 is also coupled to one or more input/output devices 710 (I/O) that include, but are not limited to, devices such as video monitors, track balls, mice, keyboards, microphones, touch-sensitive displays, transducer card readers, magnetic or paper tape readers, tablets, styluses, voice or handwriting recognizers, or other well-known input devices such as, of course, other computers. Finally, CPU 702 optionally may be coupled to a computer or telecommunications network using a network connection as shown generally at 712. The above-described devices and materials will be familiar to those of skill in the computer hardware and software arts.

As discussed in the background portion of this application, data compression techniques typically employ three steps in order to efficiently compress data. These include a transformation step, a quantization step and an encoding step. Two of the present inventors previously developed the lifting scheme described above for modeling data on complex manifolds. The

present invention seeks to provide a mechanism for compressing the types of data generated by the lifting scheme (and other similar schemes that utilize second generation wavelets). From the standpoint of compression, the lifting scheme effectively accomplishes the function of the transformation step.

5 Skipping next to encoding step of the general process for compressing data, a bit encoding technique known as the *zerotree* algorithm has been used in the past to encode classical wavelets. This zerotree algorithm method to encode wavelet coefficients was first introduced by J. M. Shapiro in "Embedded Image Coding Using Zerotrees of Wavelet Coefficients," 41(12):34453462, IEEE Trans. Signal Process, 1993. This technique has been extended recently
10 by A. Said and W. A. Pearlman in the manuscript "A New Fast and Efficient Image Codec Based on Set Partitioning in Hierarchical Trees", submitted to the IEEE Transactions on Circuits and Systems for Video Technology. A further explanation of this technique can also be found in "Image Compression Using the Spatial-Orientation Tree", A. Said and W. Pearlman, IEEE Report 0-7803-1254-6/93, 1993. Each of these publications is incorporated herein by reference.

15 Said-Pearlman have proposed a very fast and efficient method to bit encode wavelet coefficients resulting from a conventional wavelet transform. This method and its limitations can best be understood by referring to Figures 1a - 1d. Figure 1a is a symbolic representation of an image 100 that has undergone a wavelet transform. Typically, in a wavelet transform an image such as the square 102 is subdivided iteratively and at each level a coefficient is associated with a
20 level of the subdivision. For example, Figure 1a shows at the coarsest level a square 102 that has had a predicted coefficient value A 103 associated with the whole face of the square 102. Further iterations of the wavelet transform will further subdivide square 102 into smaller and smaller sub-squares until the pixel level is reached. At each level coefficients will be predicted to correspond to the divided sub-squares.

25 Figure 1b shows the next level of iteration of a wavelet transform upon the image 100. In this figure, square 102 has been further subdivided into four equally sized sub-squares. Note that the coefficient A 103 is still associated with the complete face of the square 102 as indicated by its presence in the center of square 102. At this level of iteration, square 102 has been further subdivided into squares 104, 106, 108 and 110. Using the wavelet transform and previously
30 calculated coefficients, new coefficients B 105, C 107, D 109 and E 111 are now associated with the faces of the newly divided sub-squares 104, 106, 108 and 110 respectively. In other words, coefficient A 103 is a predicted coefficient value for the face of square 102 while coefficient B 105 is a predicted value for the square 104. A predicted value might represent the gray scale of a particular square, or at the lowest level, a pixel.

35 Figure 1c represents the next level of iteration upon the image 100 using a wavelet transform. Figure 1c demonstrates how the square 102 that had been subdivided in Figure 1b is now further subdivided. For example, square 104 has now been subdivided into four sub-

squares. These sub-squares are numbered 112, 114, 116 and 118. Associated with the face of each of these smaller sub-squares is a new predicted coefficient, namely F 113, G 115, H 117, and I 119. For example, if square 112 is at the pixel level then it can be said that wavelet coefficient F 113 is associated with the whole face of square 112. Likewise, wavelet coefficient B 105 is associated with the face of square 104 and wavelet coefficient A 103 is associated with the face of square 102. This subdivision process would also be repeated for squares 106, 108 and 110, although not shown.

Figure 1d shows a tree structure 120 that represents an implementation of the zerotree method described by Said-Pearlman. This tree structure is also referred to as a quad tree or Q-Tree in that the nodes represent squares. The zerotree method first arranges all wavelet coefficients in the Q-Tree. At the root of the Q-Tree are the coefficients on the coarsest level. (Those coefficients at the top level are also referred to as the scale coefficients.) Thus, in our example, coefficient A 103 represents the coarsest level for the square 102, hence, coefficient A 103 is at the root of the Q-Tree 120. Each coefficient would then have four "children", namely the coefficients of the next finer scale of resolution that live at the same location. Thus, the children of node A will be the coefficients B 105, C 107, D 109 and E 111. In other words, the children of node A correspond to the next finer subdivision of square 102. As square 102 was subdivided into squares 104, 106, 108 and 110, the coefficients that are associated with each of these subdivided squares form the children nodes of parent node A. Nodes C, D, E would also have children corresponding to a finer subdivision (as for node B), although these children are not shown.

The encoding technique of Said-Pearlman now relies on the fact that in smooth regions of images the coefficients decay exponentially. This implies that if a certain coefficient is below a threshold then its children are very likely to be below the threshold as well. Thus, in this likely case the whole tree of small coefficients may be discarded. For example, if coefficient B is below threshold or close to zero, it is very likely that its children coefficients F, G, H and I will also be below threshold as well.

The Q-Tree 120 of Figure 1d also implicitly encodes the physical location of the coefficients. For example, because coefficients B, C, D and E are children of coefficient A, it is implicitly known that coefficient B will occupy the top left corner of square 102, coefficient C will occupy the top right corner of square 102, and so on. The zerotree encoding algorithm thus enables the coefficients to be compressed and transmitted without having to also encode each individual coordinate for each coefficient.

However, there are some limitations to this zerotree algorithm if a geometry other than square subdivision is used. For example, if a subdivision scheme based upon triangles is used and it is desired to associate wavelet coefficients with the vertices of triangles, certain problems are encountered. As described above, it is often desirable to use a triangular subdivision scheme

when modeling complex geometries. These problems will now be discussed in reference to Figures 2a, 2b and 2c.

Figure 2a shows a sphere 200 that is the subject of data compression using a wavelet transform. For example, sphere 200 might represent the Earth, another planet or a human head. In this figure the surface of sphere 200 has been subdivided into triangles in the form a geodesic sphere. Other types of subdivision with or without triangles are also possible. Shown is a triangle 210 that represents one of the many triangles on the surface of the sphere. Note that triangle 210 itself forms a part of a larger triangle that is represented by triangles 211, 210, 212 and 213. It should also be noted that triangle 210 itself may be further subdivided according to the same scheme.

Figure 2b represents a wavelet transformation process in which large triangle 210 is further subdivided into smaller triangles 212, 214, 216 and 218. Because triangle 210 is subdivided into four smaller triangles, the Q-Tree representation scheme of Figure 1d may be applicable. For example, a wavelet coefficient representing the face of triangle 210 would be at the root of the tree and the four wavelet coefficients corresponding to the faces of smaller triangles 212, 214, 216 and 218 would be the children of this parent wavelet coefficient. This would be a straightforward representation of triangle 210 if the coefficients are associated with the faces of the triangles. However, wavelet coefficients are not necessarily always associated with the faces of shapes such as triangles or squares.

In certain subdivision schemes (for example, subdivision using triangles) it may be desirable to associate the wavelet coefficients not with the face of a given unit but with its vertices. For example, in Figure 2b, triangle 210 has wavelet coefficients defined at its vertices, namely wavelet coefficients M 250, N 252, and O 254. Thus, triangle 210 and the wavelet coefficients M, N and O represent a first level in a wavelet transform. The next or second level in a wavelet transform is represented by the subdivision of triangle 210 into smaller triangles 212, 214, 216 and 218. At this next level, wavelet coefficients M, N and O may be used to predict the second level coefficients P 256, Q 258 and R 260.

Turning now to Figure 2c, a further subdivision of triangle 210 and a limitation of the zerotree algorithm is shown. In this figure, triangle 210 has been subdivided again. Top triangle 212 has now been subdivided into four smaller triangles 220, 222, 224 and 226. Wavelet coefficients may now be predicted for the new vertices. As wavelet coefficients M, N and O had been used to predict the wavelet coefficients P, Q and R, further prediction occurs similarly. For example, wavelet coefficients M, P and Q have been used to predict wavelet coefficients S 262, T 264 and U 266. Note how the wavelet coefficients S, T and U reside at the vertices of triangle 224.

It is here that a limitation of the zerotree algorithm becomes apparent. For example, if triangle 212 is considered to be a node in the Q-Tree it is unclear which wavelet coefficients should be associated with this node. And it is also unclear which wavelet coefficients should be associated with the children of triangle 212. For example, it can be seen that triangle 212 may conceivably be associated with any of the wavelet coefficients M, S, T, P, U or Q. Further compounding the problem is the fact that triangle 210 cannot be viewed in isolation. For example, referring back to Figure 2a it can be seen that the topmost wavelet coefficient M of triangle 210 is also shared with triangles 212, 211 and three other triangles. Also, if triangle 220 is one of the four children in the Q-Tree of triangle 212, it is unclear which wavelet coefficients to associate with triangle 220. Wavelet coefficients M, S and T are possibilities, but coefficients S and T reside at different levels in the subdivision scheme than coefficient M. That is, wavelet coefficient M was produced at a first iteration whereas wavelet coefficient S and T were produced at the third iteration. This discrepancy can cause problems in the zerotree algorithm. The problem is that given a vertex wavelet coefficient in the triangle tree there is no way to uniquely assign four children coefficients. This problem is caused by the fact that the wavelet coefficients reside at the vertices of the mesh and not at the faces. Note how this problem does not manifest itself if wavelet coefficients are associated with the faces of a subdivision. Referring back to Figure 1c, note that square 104 is uniquely associated with wavelet coefficient B and its subdivided squares are uniquely associated with coefficients F, G, H and I and so on. Another problem is that some wavelet coefficients have no children but do have grandchildren. Such a configuration prevents application of the original zerotree algorithm since there is no direct link between a coefficient and its grandchildren.

Referring next to Figures 3a and 3b, a method of compressing and uncompressing data in accordance with one embodiment of the invention will be described. The compression method of Figure 3a will be described first. The procedure begins at step 301. In this step, the user is prompted to input various parameters for use by the procedure. These parameters will be described more fully below with reference to steps 303 to 315. Typically, these parameters, or commands, are input by the user via a command line that is then parsed by the procedure, as in the UNIX operating system. Alternatively, these parameters may be specified by the user via a graphical user interface (GUI), or the parameters may be input via a script file or the like without user interaction. In one embodiment, the invention may be implemented in C++ on a UNIX platform. Graphical visualization may be done, for example, by using the GL library on an SGI workstation.

It is possible for the procedure to operate in an interactive mode. If interactive mode is selected the user is able to specify parameters to control. In the interactive mode, the user would be able to interact with the procedure as subdivision is occurring and the function is being compressed. For example, in the interactive mode the user would be able to perform subdivision iteratively, viewing the results after each iteration, and would also be able to rotate and view the

object from a different angle. Other options available to the user in interactive mode are the ability to view the tree structure being formed as a result of the wavelet transformation. If the interactive mode is selected, these options are available to the user as the procedure progresses. The steps of the procedure themselves are not changed, only the way in which the user can view the results.

5 In step 303 the data file is identified. The data file contains the data for a function defined upon a manifold. It is this function that will be compressed. The data file also contains a representation of the two-dimensional (2-D) manifold itself. The manifold that the function is defined on may be a simple rectangle embedded in the plane, as in a normal photograph, or it may be defined upon a more complex 2-D manifold, such as the spherical surface of the Earth,
10 which cannot be embedded in the plane. The function may be defined in any of a number of ways, including the interpolation of photographic or other sensed data. Techniques for this are well known in the art. Intuitively, one may think of the function as being "glued" onto the manifold.

15 In step 305 the level of subdivision desired is identified. The level of subdivision refers to how many times a triangle in the geometric base type is subdivided into smaller triangles. For example, with reference to Figure 2a, the sphere 200 has been subdivided at the first level. A second level of subdivision would mean subdividing each of the triangles shown on the sphere 200 into four smaller triangles, as shown in Figure 2b. The level of subdivision affects the quality of the resulting uncompressed image. For example, a high level of subdivision may
20 increase the quality of the image, but will take longer to process and requires more memory in a computer.

25 In step 307 the compression factor is identified. The compression factor is the factor by which the data representing a manifold is compressed. For example, a compression factor of four (4) would compress the data by four times (4x). In other words, a data file of 4 Megabytes would be compressed into approximately 1 Megabyte using a compression factor of four (4). A high compression factor produces a smaller compressed data file but may result in a lower quality uncompressed image.

30 In step 309 the number of bit planes is identified. The concept of a bit plane is known in the art. For example, for a graphic image, it is common to represent each value, or pixel, by 8 bits. Thus, there would be 8 bit planes, numbered zero through seven. The i-th bit plane would contain the i-th bits for each pixel in the image. In another example, an X-ray image may be represented by 12 bits for each value, resulting in 12 bit planes.

35 In step 311 the geometric base type used for modeling the manifold is identified. It is desirable to choose a geometric base type that approximates the general shape of the object whose

image data is being compressed. For example, in compressing a function defined on the surface of the Earth, an icosahedron as the geometric base type works well to approximate the spherical surface of the Earth. For a non-uniform shape, such as an image of a cat, a Dyn sphere works well. Dyn spheres are known in the art and are well suited for "blobby" type objects. For a photograph, a rectangle in the plane as the geometric base type would work best. Any geometric shape as a base type is possible. It would also be possible to define a new geometric type in order to model a unique shape.

In step 313 the function defined on the manifold is identified. The data representation of the function will be compressed. For example, the function may represent elevation, color or gray scale, or any other image characteristics. For an image of the Earth, the function may represent the elevation at a given point on the Earth. In other words, this function may be thought of as both a surface, and as a scalar valued function giving height (depth) for each point on a sphere. Any arbitrary mathematical function is also possible.

In step 315 the linear algebraic basis used for the lifting scheme is identified. A basis may be defined as a method for generating all elements as a linear combination of the basis elements. In the context of second generation wavelets, a basis defines a method for generating new coefficients from previously calculated coefficients. This basis will be used by the lifting scheme (explained more fully below) for a prediction scheme. For example, the basis may be linear, quadratic, cubic or butterfly (Dyn). It is preferable to use a Dyn basis for compressing data representing an image of the Earth. The use by the lifting scheme of a particular basis is explained in the reference "Spherical Wavelets: Efficiently Representing Functions on the Sphere", by P. Schröder, W. Sweldens and presented to SIGGRAPH on August 8, 1995, and is hereby incorporated by reference in its entirety. The use of spherical wavelets for use in the processing of spherical texture maps is described in the reference "Spherical Wavelets: Texture Processing", by P. Schröder, W. Sweldens, Rendering Techniques '95, August 1995, published by Springer, and is hereby incorporated by reference in its entirety.

In step 317 a T-Tree is created. This T-Tree (tree of triangles) contains nodes that will represent subdivided triangles of the geometric base type. These subdivided triangles are created in the data analysis step 320 below. The T-Tree is similar to the Q-Tree of Figure 1d that represents subdivided squares, except that the T-Tree is a tree representing triangles. As each triangle is subdivided into four smaller triangles (as shown in Figure 6, to be explained below), each node in the T-Tree (except for leaf nodes) will have four children nodes that represent the smaller triangles. As such, each node in the T-Tree represents a triangle in the triangular subdivision scheme. A wide variety of methods are possible for creating this tree structure. By way of example, the triangular subdivision scheme shown in the references "Spherical Wavelets: Efficiently Representing Functions on the Sphere" or "Spherical Wavelets: Texture Processing"

(both referenced above) may be used to perform subdivision, and the T-Tree will represent this subdivision. A pointer from each node in the T-Tree may be used to reference a data structure that contains information about the represented triangle including its associated wavelet coefficients.

5 In step 319 the error norms for the original data are determined. The original data is the image data before transformation or compression. Error norms are used to help evaluate the quality of the compression. For example, the maximum and the minimum of the pixel values are calculated to evaluate the quality. Other error norms calculated measure the difference between the original image and the resulting image. Mathematically, this means establishing a
10 metric for measuring distances between images. Metrics for measuring these error norms are known in the art.

In step 320 data analysis is performed to calculate the wavelet coefficients. This is a data transformation step in which the original data is transformed into a second generation wavelet representation, the data then being represented using second generation wavelet coefficients. This
15 step performs the subdivision desired upon the manifold and calculates the wavelet coefficients. A method for performing this transformation is fully described in the reference "Spherical Wavelets: Efficiently Representing Functions on the Sphere", noted above.

In step 323 the data characteristics for the uncompressed data are calculated. These characteristics include the error norms and are used for the purposes of peak signal to noise ratio
20 computation. The calculation for the peak signal to noise ratio may be performed in any manner as is known in the art.

In step 326 the compression function is executed. This function uses a tree based scheme that is somewhat analogous to that disclosed in Said-Pearlman noted above, and is explained more fully below with reference to Figures 4a, 4b and 4c. The compression function will
25 compress the wavelet coefficients representing the function defined upon the 2-D manifold. At this point, if a compression file is being written to, the file size before and after compression may be compared and displayed to the user. In step 331 the error norms for the compressed data are calculated. These are the same norms as calculated in step 319, but are calculated after compression for comparison purposes. Next, the compression procedure ends.

30 At this point, compression has been completed. If compression has been performed for a purpose such as transmitting the data, storing the data or otherwise manipulating the data, then that transmission or storage may be performed at this point and the procedure ends. In some cases, however, the user may wish to proceed directly to the uncompression procedure of Figure 3b. For example, for research, testing or other purposes, it may be desirable to uncompress the
35 data and view the resulting image immediately. In that case, the user would execute the uncompression procedure. Alternatively, if the user begins with or receives compressed data, the

user may wish to begin with the uncompression procedure and may not need to perform the compression procedure. Accordingly, the uncompression procedure of Figure 3b will now be described.

5 The uncompression procedure of Figure 3b begins at step 352. In step 354, the user identifies which data file contains the compressed data. Of course, the data may also be contained in a memory of the computer or may be received in real time over a communication line. When the data had been compressed, certain information useful to compression is saved in the compressed data file. During uncompression, this information is first read from the data file in order to indicate to the uncompression procedure how the data was compressed. This information is read in steps 356, 358, 360, 364 and 366. Of course, if compression is standardized, and this information is known a priori, the uncompression procedure would be able to make assumptions about this information and would not have to read it from a data file. For example, if a decision is made that a tetrahedron will be used for the geometric base type, then the uncompression procedure need not read this information but may assume a tetrahedron.

15 In step 356 the level of subdivision used for compression is read or received. This level corresponds to the level that was identified in step 305 of the compression procedure. In step 358 the geometric base type used for compression is read or received. This type corresponds to the type that was identified in step 311 of the compression procedure. In step 360 the basis for the lifting scheme used for compression is read or received. This basis corresponds to the basis that was identified in step 315 of the compression procedure.

In step 362 the G-Tree is created to serve as a place holder for the wavelet coefficients that will be read in from the data file using the uncompression step 368 below. This G-Tree may be created as described above in step 404 of Figure 4a, which uses the T-Tree as described in step 317 of Figure 3a.

25 In step 364 the absolute value of the maximum coefficient calculated is read or received. This value is used, for example, to convert from fixed point to floating point notation. In step 366 the number of bit planes used for compression is read or received. This number corresponds to the number that was identified in step 309 of the compression procedure. After this information has been determined, the uncompression function itself is performed.

30 In step 368 the uncompression function is executed. This step is fully explained below with reference to Figures 5a, 5b and 5c. The uncompression function will take a file of bits that had been previously written out using the compression function and uncompress this information to generate the set of wavelet coefficients that represent the function defined upon the 2-D manifold. Once the coefficients have been generated, a quality analysis may be done. Quality analysis involves calculating again the error norms for the newly uncompressed data, and

comparing to the previously calculated error norms in order to determine the quality of the compression.

In step 370 data synthesis is performed on the wavelet coefficients in order to reconstruct as best as possible the data representing the original image. This step is the reverse of step 320 of Figure 3a, calculating the wavelet coefficients. Data synthesis uses a reverse transform upon the second generation wavelet coefficients to produce the image data. A method for performing this reverse transformation is also described in the above-identified reference "Spherical Wavelets: Efficiently Representing Functions on the Sphere". At this point, the image may be displayed for evaluation.

In step 372 the user is asked if visual evaluation is desired. With an image, it is often possible to roughly evaluate the quality of compression by viewing the image before compression, and then viewing the same image after it has been compressed and then uncompressed. If visual evaluation is not desired, then the uncompression procedure ends. If visual evaluation is desired, control moves to step 374, the image is displayed and then the procedure ends.

Figures 4a, 4b and 4c are a flowchart illustrating one embodiment of the compression function step 326 of Figure 3a. The function begins in step 402. The compression function used here is analogous to Algorithm II of Said-Pearlman, although it operates upon a geometry tree (or G-Tree) instead of upon a traditional T-Tree as explained more fully below.

In step 404 the G-Tree (geometry tree) is initialized using the T-Tree that was produced by the create T-Tree step 317 of Figure 3a. The G-Tree is related to, but distinct from the T-Tree. Different schemes for forming a G-Tree are possible. One such scheme is described in detail below with reference to Figure 6. The G-Tree accommodates triangular meshes of arbitrary two-dimensional manifolds, (for example, the sphere), as contrasted with the purely rectangular system described in Said-Pearlman. In the G-Tree, the coefficients live at the vertices of triangles, and the G-Tree uniquely associates a vertex with one triangle. In the standard T-Tree, coefficients are defined at the vertices of triangles, but it is not immediately obvious how to associate a coefficient with a unique triangle.

In the standard T-Tree coefficients are defined at the vertices of triangles because the analysis step 320 that produces the second generation wavelets creates coefficients that are defined at vertices. Alternatively, however, it would be possible to generate wavelet coefficients that are defined on the faces of triangles, and a corresponding standard T-Tree. In this case, it would not be necessary to initialize a G-Tree, because a wavelet coefficient defined on the face of a triangle may uniquely be associated with that triangle.

In the current embodiment, however, the standard T-Tree represents coefficients that are defined at the vertices of triangles and it is not immediately obvious how to associate a coefficient with a unique triangle. Most vertices have six triangles incident upon them. It is necessary to associate a coefficient with a particular triangle, even though it may appear that a coefficient may be associated with more than one triangle. In other words, the G-Tree will use a tree of subdivided triangles to form its backbone, and vertices (and hence, coefficients) will be associated with a particular triangle in the G-Tree. In the G-Tree, then, a specific triangle is picked out with which to associate a vertex. There will be zero to three vertices mapped to each triangle. The following provides an explanation of how this mapping is done.

In the original Said-Pearlman reference a simple identification was made between sets of coefficients and a single coefficient in the following manner. Each coefficient has three indices, the first being the level of subdivision, where the top of the tetrahedron, i.e., the coarsest level, is numbered zero and finer levels have increasing indices. The second and third indices indicate a spatial position within a given level along the X and Y axis (in that order). The value λ is a smooth coefficient and χ^* represents any of a set of three detail coefficients (since we operate in two dimensions the detail coefficients come in sets of 3, i.e., $*$ = HL, or LH, or HH). Thus, a complete image tetrahedron contains a single $\lambda(0,0,0)$ coefficient at the top and detail coefficients at all finer levels down to some finest level $L-1$, i.e., $\chi^*(j,k,l)$ with $j=0$ to $L-1$ and k and l both = 0 to $2(j) - 1$.

Given this set of coefficients, there is a simple recursive rule by which a given coefficient stands as a representative for all coefficients living in the sub-tree underneath it:

$\lambda(0,0,0)$	enumerates	$\chi_{HL}(0,0,0), \chi_{LH}(0,0,0), \chi_{HH}(0,0,0)$
$\chi^*(j,k,l)$	enumerates	$\chi^*(j+1,2k,2l), \chi^*(j+1,2k+1,2l), \chi^*(j+1,2k,2l+1),$ $\chi^*(j+1,2k+1,2l+1)$

where "*" stands for one of HL, LH, or HH, which is arbitrary but fixed.

This enumeration scheme has to be changed to accommodate surface based compression schemes in which coefficients live at the vertices of a triangular mesh of a 2-D manifold. There is no simple rule to use a single coefficient to signify both itself and a whole set of descendent coefficients. Thus, the technique of the Said-Pearlman reference has been modified by introducing a way to represent a single coefficient as distinct from a way to represent a set of coefficients reached from a given triangle in the following manner.

A coefficient is given as a pair:

(T,E) where T is a triangle pointer and E is an edge designator.

The coefficient signified by this is the one sitting at the midpoint of the given edge. The special vertices at the top of the hierarchy (4 in the case of the tetrahedron, 8 for the octahedron, and 12 for icosahedron) are not the midpoint of any edge. They also happen to be the only sites where a smooth coefficient is located. They are thus handled separately by simply enumerating them. From then on each coefficient site, or vertex, is the midpoint of some edge.

The first difficulty arises in that a given vertex, being the midpoint of a particular edge, may have two designators:

(T, E) and (τ, E)

where T and τ are the triangles that share the common edge E . This can also be explained by noting that in working with 2-D manifolds an edge is a one-dimensional (1D) sub-manifold and each edge is incident upon either one or two triangles. Because of this a unique way is needed to resolve the ambiguity as to which is the proper designator for a given coefficient. This is done by observing that one of the triangles incident on the edge is parameterized counter-clockwise, while the other by necessity is parameterized clockwise. Because the 2-D manifold is oriented, the base geometry may be parameterized so that every pair of triangles incident upon a common edge have opposite parametrization as to clockwise versus counter-clockwise. With the subdivision and parametrization schemes described below, this property is retained at all subdivision levels.

A clockwise triangle is one in which vertices V_1 , V_2 and V_3 are enumerated in a clockwise direction. The proper designator for a given coefficient is the one for which the triangle is labeled counter-clockwise. This rule works correctly throughout the hierarchy except at the edges of the coarsest level, which are 6 for the tetrahedron, 12 for the octahedron, and 30 for the icosahedron. For these edges the ambiguity is resolved by choosing as the only legal designator amongst the two triangles sharing a given such edge, the one triangle with the lower memory address.

Sets of coefficients are now represented by a triangle alone. It represents all triangles that are below it in the hierarchy, i.e., each triangle has 4 constituent children, which themselves have possibly 4 constituent children and so forth to level $L-1$. The set of coefficients represented by some triangle T is then simply given by all coefficients that as a legal designator have:

(τ, E) for some τ a triangle descendent of the root T .

This representation allows the Said-Pearlman technique to work for data that lives on the vertices of some triangle mesh. The triangle, vertex and edge labeling scheme will now be described in detail with reference to Figure 6.

Figure 6 shows a triangle T 610 that represents one of the triangles on the surface of a sphere 200 as shown in Figure 2a, according to an embodiment of the present invention. Triangle 610 may represent a triangle at the coarsest level, or a triangle at a finer level of subdivision. The large triangle T has been subdivided into four medium triangles, each of that may be further subdivided into smaller triangles. Triangle T has vertices V1 650, V2 654 and V3 652, and has been subdivided into medium triangles T0 616 (center), T1 612 (top), T2 618 (lower right) and T3 614 (lower left). This next level of subdivision has produced vertices V1 660, V2 656 and V3 658, that are shared by new triangles T0, T1, T2 and T3.

At the next finer level of subdivision, triangle T1 is subdivided as shown (triangles T0, T2 and T3 would also be subdivided in this fashion, although not shown). Triangle T1 has vertices V1 650, V2 656 and V3 658, and has been subdivided into small triangles T1:0 624 (center), T1:1 620 (top), T1:2 626 (lower right) and T1:3 622 (lower left). This next level of subdivision has produced vertices V1 666, V2 664 and V3 662, which are shared by new triangles T1:0, T1:1, T1:2 and T1:3.

A triangle is implemented in one embodiment as a record containing four pointers to its descendent triangles (that are nil at the finest level), three pointers to vertices V1, V2, V3, a bit field indicating the parametrization of the triangle (counter-clockwise or clockwise), and a name T?, where ? = 0, 1, 2, or 3. The top level triangles are all counter-clockwise and are named T0. Thereafter, child triangles are named after the vertex that they retain from their parent and are either clockwise or counter-clockwise depending on what the order of the vertices is. For example, in Figure 6, the largest triangle is indicated by T 610, and it has child triangles T0, T1, T2 and T3. Triangle T0 has a clockwise orientation, while triangles T1, T2 and T3 all have a counter-clockwise orientation. The center triangle does not retain any of its parent vertices. An edge of a given triangle is named E? where ? = 1, 2 or 3. The number is chosen to match the vertex number of the vertex across from the given edge. As a consequence, E1 is between V2 and V3, E2 is between V3 and V1, and E3 is between V1 and V2. In Figure 6, edges E1 670, E2 674 and E3 672 are shown for triangle T0. These edges are shared respectively with triangles T1, T2 and T3. Given this, the newly introduced vertices in the middle of an edge (due to subdivision or instancing of children) are named after their parent edge. E1 generates V1, etc. For example, Figure 6 shows vertex V1 666 as the midpoint of edge E1 670.

This naming scheme is self-consistent and unambiguous. Note that triangle T0 inherits the orientation (counter-clockwise/clockwise) from its parent, while all other triangles have the reverse orientation. Accordingly, the coefficient residing at vertex V1 666 may be designated as (T1, E1) or as (T0, E1). As T1 has a counter-clockwise orientation, the correct designation is (T1, E1). This designation is arbitrary, and it is contemplated that the correct designation may be the one for which the triangle has a clockwise orientation. Thus, using the above representation,

the T-Tree may be converted into a G-Tree in which wavelet coefficients residing at vertices are associated with a unique triangle.

It is possible to utilize other schemes for forming the G-Tree. By way of example, for the subdivision scheme of Figure 6, a G-Tree may be formed where the nodes of the G-Tree are any of triangles, edges, or vertices. The base geometry forms the root of the G-Tree and immediately below it are each of the vertices, edges and triangles of the base geometry. To subdivide a triangle in this fashion, its edges are first subdivided. Each edge subdivides into two edges and a (central) vertex. Each of these is a G-Tree child of the subdivided edge. The triangle is now subdivided into four triangles and three edges (the interior edges) that become the G-Tree children of the subdivided triangle. Each vertex is now uniquely assigned to either the base geometry or to an edge. Each edge is uniquely assigned to one of the base geometry, another edge, or a triangle. Each triangle is uniquely assigned to one of the base geometry or another triangle.

The preceding scheme for subdivision may be generalized to M-dimensional manifolds. Vertices, edges, and triangles are all instances of cells (0-cells, 1-cells and 2-cells respectively). Cells have a dimension, and cells of dimension M are called M-cells. This concept of cells is well known to those of skill in the art. A generalization to M-dimensional manifolds may be performed as follows. Subdividing a cell requires a previous subdivision of the cells on its boundary (all of which are of strictly lower dimension). The subdivision of a 0-cell is itself. The cells in the subdivision become the G-Tree children of the subdivided cell, and are associated with the subdivided cell. The cells of the base geometry are the G-Tree children of the base geometry itself, and each is associated with the base geometry.

Within this generalized scheme there are a number of methods for performing the actual cell subdivision. One such subdivision scheme is *barycentric* subdivision which is known in the art. Figure 7a illustrates the barycentric subdivision of a triangle. Barycentric subdivision subdivides an M-cell into a (central) 0-cell and a K+1 cell for each K-cell of the subdivided boundary of the M-cell. Each of the K-cells in the subdivision have K+1 0-cells on their boundary and is a K-dimensional generalization of the triangle (called a *simplex*, a 2-simplex is a triangle, a 3-simplex is a pyramid). Note that the barycentric subdivision of a triangle has 13 cells (the central 0-cell, six 1-cells from the six vertices of the subdivided boundary, and six 2-cells from the six edges of the subdivided boundary).

Another scheme for the actual cell subdivision is *bi-conic* subdivision. Bi-conic subdivision of a triangle is shown in Figure 7b, and of a quadrilateral in Figure 7c. Bi-conic subdivision subdivides an M-cell into a (central) 0-cell and an M-K cell for each K-cell of the unsubdivided boundary of the M-cell. Each of the K-cells in the subdivision have at least K+1 0-cells on their boundary. A 2-cell in the subdivision will be a quadrilateral. These various subdivision schemes are illustrative of the ways in which a cell may be subdivided. Once the G-Tree has been created here in step 404, the compression function continues.

In step 406 the wavelet coefficients are converted to binary form. Specifically, each wavelet coefficient in scientific floating point notation is converted to fixed point binary signed magnitude notation. This conversion may be accomplished by walking the G-Tree and converting each coefficient in turn.

5 In step 408 the significance bits "S(N)" for each wavelet coefficient are calculated. In general, the calculation of significance bits is known in the art. This significance function is analogous to the one described in Said-Pearlman except that the tree on which the S(N) bits are calculated is not the Q-Tree of Said-Pearlman but the G-Tree. The "S" function is a mask that indicates whether a bit plane is significant for a sub-tree. More specifically, for a given wavelet
 10 coefficient of N bits, the corresponding significance function S(N) will be a "1" if a bit at bit position N is not a leading "0" (or if S(N) for a descendant coefficient is a "1"), and S(N) will be a "0" if the bit at position N is a leading zero and S(N) for every descendent coefficient is also a "0". For example, if a wavelet coefficient has a binary representation of "0001 0100", then the corresponding S(N) function would be "0001 1111", the first three binary digits of S(N) being
 15 each a "0" because the wavelet coefficient has three leading zeros.

As alluded to above, S(N) for a given wavelet coefficient in the G-Tree will also be a "1" if for that bit position N, there exists a descendant wavelet coefficient with a corresponding S(N) that also has a "1" in that position. In other words, S(N) for a given wavelet coefficient "W" in the G-Tree is the logical "or" of the S(N) functions corresponding to all wavelet coefficients that
 20 are descendants of "W" in the G-Tree. Using the above example, if "W" is a wavelet coefficient represented as "0001 0100", and there exists a descendent of "W", "X" with a coefficient represented as "0011 0000", then S(N) for "X" is "0011 1111", and S(N) for "W" would also be "0011 1111".

Steps 410 through 466 of Figures 4a, 4b and 4c implement a procedure analogous to
 25 Algorithm II of Said-Pearlman, although the data being processed is quite different since a G-Tree is used. This procedure uses various lists and sets to facilitate its implementation. LIS stands for "List of Insignificant Sets", LIN is "List of Insignificant Nodes" and LSN is "List of Significant Nodes." The sets used are O(T), all offspring or children of a node T, D(T), all descendants of a node T, and L(T) which is D(T) minus O(T) (all descendants of node T except
 30 for immediate children). This procedure will now be explained.

In step 410 the bit position N is set equal to the most significant bit of the binary representation of the wavelet coefficient. For example, if 8 bits are used to represent the wavelet coefficients, then N = 7. For the least significant bit, N = 0. This bit position N is also preferably
 35 output to a file, that will contain the compressed data when the compression step 326 of Figure 3a is finished. The other "output" steps referred to below also indicate that the bit is to be output to this same file. However, it is not necessary that the bits be output to a file. It is also contemplated that the compressed bits may be stored in memory, written to a disk, or otherwise

stored in a computer system. Furthermore, the bits need not be stored locally, but may be transmitted upon output via a computer network, telephone line, microwave, satellite transmission or the like to be processed at a remote location.

In step 412 the ordered lists LIS (List of Insignificant Sets) and LIN (List of Insignificant Nodes) are initialized to contain all of the top level nodes. A top level node represents a wavelet coefficient at the coarsest level of the subdivision, as a result of the analysis step 320 of Figure 3a. In step 414 the list LSN (List of Significant Nodes) is initialized to the empty list.

Step 416 represents a looping structure in which index T is used to step through each entry in the LIN. At first, T is initialized to the first entry in the LIN. If the end of the LIN has been reached then control transfers to step 426 of Figure 4b. If the end of the LIN has not been reached then the loop continues through steps 418 to 424.

In Step 418 the significance bit, $S(N)$, is output for node T. Because there may be from 0 to 3 wavelet coefficients associated with this node T, there will be from 0 to 3 significance bits output in this step, depending upon how many coefficients are associated with this node. Step 420 tests whether the significance of the bit at bit position N for any of the coefficients associated with node T is equal to 1. If not, control returns to step 416 and T is set equal to the next entry in the LIN. If so, step 422 is executed in which node T is moved to the LSN. Next, in step 424, the sign of each of the coefficients at node T is output. After step 424 control returns to step 416 in which T is set equal to the next entry in the LIN. As soon as the end of the LIN is reached control moves to step 426 of Figure 4b.

Step 426 of Figure 4b is a looping structure that steps through each entry in the LIS until the end of the LIS is reached. Initially, T is set equal to the first entry in the LIS. When the end of the LIS is reached control moves to step 458 of Figure 4c. Step 428 tests whether node T is marked as representing set D. If so, then steps 430-448 are executed. If not, then T is marked as representing set L and control moves to step 450 of Figure 4c where steps 450-456 are executed. In step 430 all of the significance bits at bit position N for all the coefficients in set D(T) are output. Step 432 tests whether the significance bit at bit position N for any of the coefficients in set D(T) is equal to one. If the significance bit is not one then control returns to step 426 and the counter T is set equal to the next entry in the LIS. If the significance bit is a one then control moves to step 434.

Step 434 is a looping structure in which an index U is used to step through each entry in the set O(T). Initially U is set equal to the first entry in set O(T). If the end of set O(T) is reached then control moves to step 446. In step 436 the significance bit at bit position N for each of the coefficients at node U is output. Step 438 tests whether the significance bit at bit position N for any of the coefficients associated with node U is equal to one. If not, control moves to step 444. If so, control continues to step 440. In step 440 node U is added to the LSN. In step 442 the sign

of each of the coefficients at node U is output. After step 442 control returns to step 434 and U is set equal to the next entry in set O(T). In step 444 node U is added to the LIN and then control returns to step 434. Step 446 tests whether set L(T) is empty. If so, then control returns to step 426 and T is set equal to the next entry in the LIS. If set L(T) is not empty then in step 448 node
 5 T is moved to the end of the LIS and is marked as representing set L. After step 448 control returns to step 426.

In step 450 of Figure 4c the significance bits at bit position N for each of the coefficients in the set L(T) are output. Step 452 tests whether the significance bit at bit position N for any of the coefficients associated with the set L(T) is equal to a one. If the significance bit is not a one
 10 then control returns to step 426 of Figure 4b. If the significance bit is a one then step 454 is executed. In step 454 each member of set O(T) is added to the end of the LIS and is marked as representing set D. In step 456 node T is removed from the LIS and control then returns to step 426 of Figure 4b.

Step 458 represents a looping structure in which an index T is used to step through each
 15 entry in the LSN. Initially T is set equal to the first entry in the LSN. If the end of the LSN is reached then control moves to step 464. Step 460 tests whether the entry T in LSN has been included in the most recent sorting pass through LIN and LIS. If so then control returns to step 458 and T is set equal to the next entry in LSN. If not, then step 462 is executed. In step 462 the Nth most significant bit of the absolute value of each of the coefficients at node T is output. After
 20 step 462 control returns to step 458. In step 464 the counter N is decremented by one. Step 466 checks whether N is equal to zero. As N represents the number of bit planes for the coefficients, if N is equal to zero then step 326 of Figure 3b is done. If N is not zero then control returns to step 416 of Figure 4a and T is set equal to the next entry in the LIN.

Turning now to Figures 5a, 5b and 5c, step 368 of Figure 3b "execute uncompression
 25 function" will now be explained in detail. Steps 510 through 566 of Figures 5a, 5b and 5c implement an uncompression procedure analogous to Algorithm II of Said-Pearlman, although the following procedure uncompresses data that has previously been compressed, whereas a literal reading of Said-Pearlman produces a compression algorithm. Nonetheless, the reader will appreciate that this uncompression procedure is analogous. Figure 5a shows a flowchart
 30 representing the uncompression procedure step 368 of Figure 3b. The procedure begins at step 502.

In step 510 the most significant bit N in the binary representation of the wavelet coefficients is read in from the output file. N now represents the current bit position, starting with the most significant bit. This bit position N is preferably input from a file that contains the
 35 compressed data after the compression step 326 of Figure 3a has finished. The other "input" steps referred to below also indicate that the bit is to be input from this same file. However, it is not necessary that the bits be input from a file. As described above, it is also contemplated that

the compressed bits may be read from memory, a disk, or other storage area in a computer system. Furthermore, the bits need not be read locally, but may be read from a computer network, telephone line, microwave, satellite transmission or the like and then be processed locally.

5 In step 512 the ordered lists LIS and LIN are initialized to contain all of the top level nodes. These top level nodes may be read from the output file where they were written previously in the compression step. A top level node represents a scale coefficient at the coarsest level of the subdivision, as a result of the analysis step 320 of Figure 3a. In step 514 the list LSN is initialized to the empty list.

10 Step 516 represents a looping structure in which index T is used to step through each entry in the LIN. At first, T is initialized to the first entry in the LIN. If the end of the LIN has been reached then control transfers to step 526 of Figure 5b. If the end of the LIN has not been reached then the loop continues through steps 518 to 524.

15 In Step 518 the significance bit, $S(N)$, is input for each of the coefficients associated with node T. Step 520 tests whether the significance of the bit at bit position N for any of the coefficients of node T is equal to 1. If not, control returns to step 516 and T is set equal to the next entry in the LIN. If so, step 522 is executed in which node T is moved to the LSN. Next, in step 524, the sign of each of the coefficients at node T is input. After step 524 control returns to step 516 in which T is set equal to the next entry in the LIN. As soon as the end of the LIN is reached control moves to step 526 of Figure 5b.

20 Step 526 of Figure 5b is a looping structure that steps through each entry in the LIS until the end of the LIS is reached. Initially, T is set equal to the first entry in the LIS. When the end of the LIS is reached control moves to step 558 of Figure 5c. Step 528 tests whether node T is marked as representing set D. If so, then steps 530-548 are executed. If not, then T is marked as representing set L and control moves to step 550 of Figure 5c where steps 550-556 are executed. In step 530 all of the significance bits at bit position N for all the coefficients in set D(T) are input. Step 532 tests whether the significance bit at bit position N for any of the coefficients in set D(T) is equal to one. If the significance bit is not one then control returns to step 526 and the counter T is set equal to the next entry in the LIS. If the significance bit is a one then control moves to step 30 534.

35 Step 534 is a looping structure in which an index U is used to step through each entry in the set O(T). Initially U is set equal to the first entry in set O(T). If the end of set O(T) is reached then control moves to step 546. In step 536 the significance bit at bit position N for each of the coefficients at node U is input. Step 538 tests whether the significance bit at bit position N for any of the coefficients associated with node U is equal to one. If not, control moves to step 544. If so, control continues to step 540. In step 540 node U is added to the LSN. In step 542 the sign

of each of the coefficients at node U is input. After step 542 control returns to step 534 and U is set equal to the next entry in set O(T). In step 544 node U is added to the LIN and then control returns to step 534. Step 546 tests whether set L(T) is empty. If so, then control returns to step 526 and T is set equal to the next entry in the LIS. If set L(T) is not empty then in step 548 node
5 T is moved to the end of the LIS and is marked as representing set L. After step 548 control returns to step 526.

In step 550 of Figure 5c the significance bits at bit position N for the coefficients in the set L(T) are input. Step 552 tests whether the significance bit at bit position N for any of the coefficients in the set L(T) is equal to a one. If the significance bit is not a one then control returns
10 to step 526 of Figure 5b. If the significance bit is a one then step 554 is executed. In step 554 each member of set O(T) is added to the end of the LIS and is marked as representing set D. In step 556 node T is removed from the LIS and control then returns to step 526 of Figure 5b.

Step 558 represents a looping structure in which an index T is used to step through each entry in the LSN. Initially T is set equal to the first entry in the LSN. If the end of the LSN is
15 reached then control moves to step 564. Step 560 tests whether the entry T in LSN has been included in the most recent sorting pass through LIN and LIS. If so then control returns to step 558 and T is set equal to the next entry in LSN. If not, then step 562 is executed. In step 562 the Nth most significant bit of the absolute value of each of the coefficients at node T is input. After step 562 control returns to step 558. In step 564 the counter N is decremented by one. Step 566
20 checks whether N is equal to zero. As N represents the number of bit planes for the coefficients, if N is equal to zero then step 368 of Figure 3b is done. If N is not zero then control returns to step 516 of Figure 5a and T is set equal to the next entry in the LIN.

Although the foregoing invention has been described in some detail for purposes of clarity of understanding, it will be apparent that certain changes and modifications may be practiced
25 within the scope of the appended claims. For instance, the image or manifold on which a function is to be compressed need not be a sphere, but may be any other two-dimensional manifold. Also, the geometric base type may be any geometric shape, or even a specially defined geometry. The invention may also be practiced upon an M-dimensional manifold that is embedded in N-dimensional space, where M and N are integers and N is greater than M. In
30 addition, the subdivision scheme may use any geometric shape for performing the subdivision and need not be restricted to triangles. By way of example, in higher dimensions the geometric base may be subdivided into cells. A cell is a unit of subdivision and is known in the art. For a 2-D manifold the cell may be any polygon, and for a 3-D manifold the cell may be any polytope. And although a second generation wavelet transform is described to perform the data analysis and
35 generate the wavelet coefficients, it is contemplated that other data transformation techniques may be used to generate the coefficients and build up a tree structure.

And within the G-Tree, although only one technique has been described for uniquely associating a coefficient defined at a vertex with a triangle, it may be possible to use other techniques for associating a coefficient with one unique triangle. Additionally, a coefficient may be associated with the face of a triangle, meaning that the coefficient would be uniquely associated with that triangle. Also, although an embodiment has been described that operates upon a static data set defined upon a sphere, for example, the invention may also be used to compress and uncompress dynamically-changing, time-dependent data sets. Such data sets occur in applications such as animation, modeling of natural phenomena, etc.

EXEMPLARY EMBODIMENT

In one embodiment of the invention, optimization of the present invention for high-dimensional compression is achieved through combining discrete second generation wavelet transforms with an extension of the embedded zerotree coding method. Other optimizations include the use of arithmetic coding, different scaling and norms of the wavelet coefficients. Also, different prediction schemes in the zerotree mechanism may be used. Combining these techniques can lead to a significant improvement in compression performance. A discussion of other embodiments are described in "Compression of Functions Defined on Surfaces of 3D Objects", K. Kolarov and W. Lynch, Proceedings of the Data Compression Conference (1997), which is hereby incorporated by reference for jurisdictions permitting such incorporation.

This embodiment discusses optimizations through the use of (1) Scaling, (2) Parameter Variation, and (3) Entropy Coding which are discussed below.

Example data sets used for optimization are an Earth image data set and the Lena image. For the Earth example a function that is the elevation (with respect to sea level) of the Earth is compressed. This topographic function is approximated from the ETOPO10 data set which samples the Earth from satellites every 10 arc minutes (1.5 million points approximately 11 miles apart). For the base complex, an icosahedron is used and is subdivided eight times, yielding 655362 vertices and 1310720 triangles. The vertices of this complex are radially projected onto a concentric geoid (vertices average about a 22 arc minute separation - about 25 miles). The topographic elevation at each vertex is then determined by interpolation of the ETOPO10 data set (2:1 reduction in points).

For comparison with existing data compression techniques the Lena image is compressed using an embodiment of the present invention (the function being compressed is the gray-scale level at each pixel). A base complex includes two triangles, thus forming a square. This complex is subdivided nine times yielding a complex with 513^2 vertices and $2 * 512^2$ triangles. The function value at each of the 513^2 vertices is determined by interpolation from a 512^2 gray-scale Lena image.

The evaluation of a compression technique may be performed in a wide variety of ways, and may depend upon the compression technique used to return a meaningful result. By way of example, for the present invention, evaluation of compression by looking at the amount of information that is read during decompression in relation to the amount of information that gets written out by the compression module works well. Parameters that may be analyzed to determine compression performance include bits per pixel, a PSNR, and a compression ratio.

One of the input parameters to an embodiment of the present invention is the number of bit planes used for representing each wavelet coefficient. Variation of the number of bit planes can affect compression performance. During compression the wavelet coefficients generated by the lifting are scaled and then further processed by the zerotree construction. The resulting output stream consists of coefficients (in order of magnitude) from the most significant to the least significant bit planes. This output stream includes all of the bit planes for the coefficients as truncated by the zerotree hypothesis. At decompression it can be decided (based on time and space resources) to read in all or just a portion of the output bit stream written during compression. This ability reflects the "progressive coding" character of the code.

A value for bits per pixel is a good measure of compression performance which together with the PSNR can be used for comparison with other algorithms. The notion of bits per pixel corresponds to bits per (wavelet) coefficient, since the coefficients are the ones that represent the data structure. Using the Lena image as an example, for 10 bit planes read in during decompression, the actual size of the file is 369189 bits for the Lena image. Those bits correspond to $513^2 = 263169$ wavelet coefficients, yielding a rate of 1.4029 bits per pixel.

A compression ratio is a widely used characteristic for comparison of algorithms. The compression ratio is calculated with respect to the amount of data needed to write out directly all of the wavelet coefficients. Each coefficient takes 8 bits for complete representation. Thus, in the example of the Lena image, 263169 coefficients will need 2105352 bits for complete representation. Using 10 bit planes per coefficient the present invention reads (and first writes) 369189 bits as outlined. Thus, the true compression ratio is 6:1 for these parameters. A 1:1 ratio can be achieved by taking more bit planes per coefficient, for example, 15 bit planes achieves that ratio.

SCALING

During decompression the wavelets are descaled by the inverse of the scaling factor used during compression. As explained above, optimization for the different norms can be achieved by scaling the coefficients after the transform analysis and before the zerotree compression for a given threshold. Multiplication by unity optimizes for the L_∞ norm. To optimize for the L_2 norm the wavelet coefficients may be scaled by 2^{-j} where j is the subdivision level at which the coefficient is attached. This basically corresponds to shifting the bits in the coefficients'

representation one to the left. To optimize for the L_1 norm, the wavelet coefficients may be scaled by 4^{-j} . This corresponds to two bits shift in the coefficients. For the general L_p norm the scaling factor is $4^{\frac{-j}{p}}$. The contrast sensitivity curve for the human visual system is best approximated by using the L_1 norm and its associated scaling.

5 Comparative results for the Lena image with L_1 and L_∞ reveal that the performance curves for the L_1 and the L_2 scaling are actually the same. All points using the L_1 norm lie on the curve for L_1 but are much earlier on the curve than the corresponding bit planes points from L_2 norm. L_1 points further along the curve can be obtained by using more bit planes for the coefficients (e.g. 15 bit planes yield a point with 32.36 dB for 0.45 b/vertex). Since L_1 vs. L_2 involves shifting two positions left for L_1 vs. one position left for L_2 , it is natural to expect that a smaller output stream results, and yields a smaller bits per pixel rate and a correspondingly smaller PSNR for the same number of bit planes read. Also, any scaling of L_1 or L_2 performs better than no scaling at all (i.e. L_∞ norm). Using scaling results in a 6 dB improvement over the whole range of the performance curve. L_2 scaling of coefficients is used for the below variations in parameters and coding for this embodiment.

PARAMETER VARIATION

Varying the different parameters used in the technique of the present invention affects compression performance. This section considers the influence of the number of levels of subdivision used, the role of the number of bit planes used to represent the coefficients, and the use of different prediction schemes.

The levels of subdivision of the base manifold affects compression performance. There is an advantage to using as many subdivision levels as possible in order to cover all of the data points in the original database. Adding one more level of subdivision increases the PSNR with more than 5 dB for both the Lena image and the Earth data. Thus, as much subdivision as possible is advantageous, although it may be constrained by hardware limitations.

The number of bit planes to represent the wavelet coefficients also affects compression performance. A wide variety of number of bit planes may be used. Ten bit planes works well, although in some cases it may be desirable to have more bit planes in order to obtain a higher PSNR and better visual perception quality. Naturally, for images like Lena the maximal bits per pixel (coefficient) rate that is needed is 8 bits per pixel, since that is what it takes to represent all coefficients explicitly. To obtain a 1:1 basis for visual comparison of images, it is useful to determine the maximal number of bit planes needed.

For example, the Lena image using 9 levels of subdivision can obtain a 1:1 compression ratio with a 15 bit plane representation of the coefficients. The resulting PSNR is 66.56 dB for 8.15 bits per pixel. Calculating the performance curve for reading in at decompression using 15,

14, 13, ..., 2, 1 bit planes of the 15 bit planes written out at compression yields the same curve with the new points beyond 10 bit planes lying further to the right on the curve.

Using a linear vs. a butterfly prediction scheme may also affects compression performance. The above examples use the linear prediction scheme as outlined above. In the linear prediction scheme, the new wavelet coefficients generated at the midpoints of any given edge are calculated as the mean of the two coefficients associated with the vertices of the edge. Another scheme is the butterfly scheme where the new coefficient is a linear combination of eight neighboring vertices from the previous level, i.e. a larger stencil is used for prediction of new coefficients. Using the Earth data example, 8 levels of subdivision and butterfly prediction brings about a 4% improvement over the linear case. The geometry of the mesh approximation of the Earth sphere being quite regular may be responsible for this slight increase. Improvements in performance may be seen when using larger stencils for high-dimensional objects with sharp edges and concavities.

ENTROPY CODING

The addition of an entropy coder and modifications to the coding scheme also affect the performance of the present invention.

Arithmetic coding of the output binary stream can be advantageous. The result from applying a wavelet transform, scaling of the wavelet coefficients and the zerotree compression is an output bit stream. The performance curve can be further enhanced by applying Huffman coding or arithmetic coding to this output stream.

A wide variety of arithmetic coders may be used. By way of example, the well-known Witten coder may be used. In one embodiment, the present invention generates the arithmetic coded output stream during decompression in parallel with reading in the compressed data. The PSNR is not affected by the arithmetic coding; the size of the bit stream is smaller. Using the Lena image as an example, using 9 levels of subdivision reveals that the arithmetic coder brings on the average 3.3% improvement over straight binary coding.

An alternative way of building the G-tree in the zerotree implementation helps to achieve a significant improvement of more than 12% compared to the results without this alternative way. One implementation of the zerotree compression using G-trees is described in detail above. The main structure of that implementation bears similarity with the SPIHT (Algorithm II) described in "A New Fast and Efficient Image Codec Based on Set Partitioning in Hierarchical Trees" referenced above. This implementation uses new data structures (the G-tree) and control flow to account for the more complicated topology of the spaces to be processed. Varying the initialization and the branching phases affects compression performance as will now be described.

G IT is defined as a list of triangles in the G-tree, GIV and GSV are lists of vertices in the G-tree. The first part of the algorithm for initialization is:

Set $n = N = \text{sig}(D(R))$ output N .

Initialize GIT to the triangles in the base complex (as type A triangles).

5 Initialize GIV to the vertices in the base complex.

Initialize GSV to empty.

Here $D(R)$ is the set of descendants of the root of the tree R .

In other embodiments, the top level coefficients at the vertices (4 for the square defining the Lena image and 20 for the icosahedron defining the Earth approximation) may be written out
10 directly in the output stream. This creates a choice for the initialization of the data structures.

A first option is to initialize GIT with the top level triangular faces (2 for the square and 20 for the icosahedron) and initialize GIV to empty (since the top level vertices are not processed). Using this option the fan out on the next level per vertex is maximum 4 and good depth can be achieved in the tree. The significance of a vertex is the value that is calculated for it. The
15 significance of a triangle is based on the significance of the triangles into which it is subdivided as well as the significance of its vertices.

A second option is to initialize GIV with the midpoints of the edges of the base triangles in GIT (5 in the case of the square and 30 for the icosahedron). The original top level vertices are not used because they are not processed. This structures calculates significance of all triangles of
20 Type A by looking at their children (triangles) in the tree and their vertices. The significance of all triangles of type B (already encountered) is calculated using the significance of the descendants except itself and its children. The fan out of this approach is 8 per node and the G tree is generally not as deep as in the previous case.

The first option has very similar performance curve to the second option. The
25 improvement for the second option is about 4% on the average for the Lena image and 1.75% on the average for the Earth. However if arithmetic coding is added to the wider G-tree generation a significant improvement is achieved. The compression curve becomes on the average 10% better than the resolution. The improvement in performance is more than 1 dB.

Thus, the butterfly scheme allows a 4% improvement over the linear one. For a span of 4
30 children per vertex the arithmetic coding improves the performance by 4% on the average. For the span of 8 children the arithmetic coder brings about 10% improvement. Good results are achieved when as many subdivision levels as necessary are used to cover the original data. More bit planes allow compression ratios closer to 1:1 and help to obtain base pictures for comparison.

Also, it is beneficial to use scaling and L_2 works well as an appropriate norm for a good compression curve. Using a combination of these techniques brings about 12% improvement in performance with respect to the results without these techniques.

5 The invention has applications in telecommunications, on the Internet, in database and computer graphic applications and in similar industries. Any application where complex data is being manipulated, compressed or transmitted may benefit from use of the invention. Therefore, the described embodiment should be taken as illustrative and not restrictive, and the invention should be limited solely by the scope of the following claims.

CLAIMS

1. A computer-implemented method for compressing data representing a function defined upon an M-dimensional manifold, the manifold being embedded in N-dimensional space where M and N are integers and N is greater than M, a geometric base being used as a coarse
5 initial model of the manifold, the method comprising the steps of:

calculating second generation wavelets for the function using a subdivision scheme for subdividing the geometric base into cells in order to produce a refined mesh, the wavelet coefficients being defined on cells within the refined mesh;

10 defining a tree structure in which each node of the tree structure represents an associated cell of the refined mesh, there being a plurality of levels in the tree, the tree levels including a top level that includes at least one top level node, wherein each tree level below the top level includes a plurality of nodes that are children of parent nodes in a next higher level, wherein children nodes associated with a particular parent node represent cells of the subdivided cell associated with the particular parent node, wherein each wavelet coefficient at a particular level in the refined mesh is
15 uniquely defined on a particular cell in the refined mesh at that level and is uniquely assigned to a single one of the cells, such that each cell may have assigned to it any number of wavelet coefficients;

setting a first threshold and outputting a stream of bits indicative of significant nodes, wherein a significant node is defined as a node that has at least one of, associated wavelet
20 coefficient that exceeds the designated threshold and descendant node that has an associated wavelet coefficient that exceeds the designated threshold, wherein when a particular significant node is determined, the output stream of bits further includes bits indicative of the significance of children nodes of the particular significant node and the significant children are recursively examined for significant descendants and the output bit stream identifies the significant
25 descendants.

2. A method as recited in claim 1 wherein each time a significant node is identified, an indicia of the sign of each significant coefficient associated with the significant node is also output.

3. A method as recited in claim 1 further comprising a refinement step that refines
30 the setting and outputting step, the refinement step including the substeps of:

setting a current threshold to a next threshold;

outputting a stream of bits indicative of newly significant nodes at the current threshold, wherein a newly significant node is defined as a node that is significant at the current threshold but was not significant at any previous threshold; and

5 each time a new coefficient that exceeds the current threshold is identified, outputting an indicia of the sign of the new coefficient.

4. A method as recited in claim 3 further comprising the step of repeating the refinement step wherein the current threshold is set to a threshold that is one-half of the previous threshold.

10 5. A method as recited in any of claims 1-4 wherein the M-dimensional manifold embedded in N-dimensional space is a 2-dimensional manifold embedded in 3-dimensional space.

15 6. A method as recited in any of claims 1-5 wherein the subdivision scheme is a triangular subdivision scheme, the cells of the subdivided geometric base are triangles and the refined mesh is a refined triangular mesh that contemplates that each base triangle in a selected level of the triangular mesh is divided into four subtriangles at the next lower level of subdivision, such that the vertices of the subtriangles include the vertices of the base triangle and the midpoints of the edges of the base triangle.

20 7. A method as recited in any of claims 1-6 wherein the calculating step uses a lifting scheme with a basis that is selected from the group consisting of linear, quadratic, cubic and butterfly.

25 8. A method as recited in claim 6 wherein each wavelet coefficient is defined at the vertex of a triangle within the refined triangular mesh and the particular vertex at which each wavelet coefficient is defined is incident upon the midpoint of an edge belonging to two larger triangles at a next higher level of subdivision, and the wavelet coefficient is uniquely associated with one of the two larger triangles.

30 9. A computer program product comprising a computer-usable medium having computer-readable code embodied thereon for compressing data representing a function defined upon an M-dimensional manifold, the manifold being embedded in N-dimensional space where M and N are integers and N is greater than M, a geometric base being used as a coarse initial model of the manifold, the computer program product comprising the following computer-readable program code for effecting actions in the computer:

program code for calculating second generation wavelets for the function using a subdivision scheme for subdividing the geometric base into cells in order to produce a refined mesh, the wavelet coefficients being defined on cells within the refined mesh;

program code for defining a tree structure in which each node of the tree structure represents an associated cell of the refined mesh, there being a plurality of levels in the tree, the tree levels including a top level that includes at least one top level node, wherein each tree level below the top level includes a plurality of nodes that are children of parent nodes in a next higher level, wherein children nodes associated with a particular parent node represent cells of the subdivided cell associated with the particular parent node, wherein each wavelet coefficient at a particular level in the refined mesh is uniquely defined on a particular cell in the refined mesh at that level and is uniquely assigned to a single one of the cells, such that each cell may have assigned to it any number of wavelet coefficients;

program code for setting a first threshold and outputting a stream of bits indicative of significant nodes, wherein a significant node is defined as a node that has at least one of, associated wavelet coefficient that exceeds the designated threshold and descendant node that has an associated wavelet coefficient that exceeds the designated threshold, wherein when a particular significant node is determined, the output stream of bits further includes bits indicative of the significance of children nodes of the particular significant node and the significant children are recursively examined for significant descendants and the output bit stream identifies the significant descendants.

10. A computer-implemented method for compressing data representing a function defined upon a two-dimensional manifold, the manifold being embedded in three-dimensional space, a geometric base being used as a coarse initial model of the manifold, the method comprising the steps of:

calculating second generation wavelets for the function using a triangular subdivision scheme for subdividing the geometric base in order to produce a refined triangular mesh, the wavelet coefficients being defined at the vertices of the triangles in the triangular mesh;

defining a tree structure in which each node of the tree structure represents an associated triangle of the triangular mesh, there being a plurality of levels in the tree, the tree levels including a top level that includes at least one top level node, wherein each tree level below the top level includes a plurality of nodes that are children of parent nodes in a next higher level, wherein children nodes associated with a particular parent node represent triangles of the subdivided triangle associated with the particular parent node, wherein each wavelet coefficient at a particular level in the triangular mesh is uniquely defined at a particular vertex in the triangular mesh at that level and is uniquely assigned to a single one of the triangles at a next higher level of subdivision, such that each triangle at the next higher level of subdivision has in the range of zero to three assigned wavelet coefficients;

setting a first threshold and outputting a stream of bits indicative of significant nodes, wherein a significant node is defined as a node that has at least one of, associated wavelet

coefficient that exceeds the designated threshold and descendant node that has an associated wavelet coefficient that exceeds the designated threshold, wherein when a particular significant node is determined, the output stream of bits further includes bits indicative of the significance of children nodes of the particular significant node and the significant children are recursively examined for significant descendants and the output bit stream identifies the significant descendants.

11. A method as recited in claim 10 wherein each time a significant node is identified, an indicia of the sign of each significant coefficient associated with the significant node is also output.

10 12. A method as recited in claim 10 further comprising a refinement step that refines the setting and outputting step, the refinement step including the substeps of:

setting a current threshold to a next threshold;

15 outputting a stream of bits indicative of newly significant nodes at the current threshold, wherein a newly significant node is defined as a node that is significant at the current threshold but was not significant at any previous threshold; and

each time a new coefficient that exceeds the current threshold is identified, outputting an indicia of the sign of the new coefficient.

20 13. A method as recited in claim 12 further comprising the step of repeating the refinement step wherein the current threshold is set to a threshold that is one-half of the previous threshold.

14. A method as recited in any of claims 10-13 wherein the manifold is a sphere and the geometric base is selected from the group consisting of an icosahedron, an octahedron and a tetrahedron.

25 15. A method as recited in any of claims 10-14 wherein the triangular subdivision scheme contemplates that each base triangle in a selected level of the mesh is divided into four subtriangles at the next lower level of subdivision, such that the vertices of the subtriangles include the vertices of the base triangle and the midpoints of the edges of the base triangle.

30 16. A method as recited in any of claims 10-15 wherein the calculating step uses a lifting scheme with a basis that is selected from the group consisting of linear, quadratic, cubic and butterfly.

17. A method as recited in any of claims 10-16 wherein the particular vertex at which each wavelet coefficient is defined is incident upon the midpoint of an edge belonging to two

larger triangles at a next higher level of subdivision, and the wavelet coefficient is uniquely associated with one of the two larger triangles.

18. A computer program product comprising a computer-usable medium having computer-readable code embodied thereon for compressing data representing a function defined upon a two-dimensional manifold, the manifold being embedded in three-dimensional space, a geometric base being used as a coarse initial model of the manifold, the computer program product comprising the following computer-readable program code for effecting actions in the computer:

program code for calculating second generation wavelets for the function using a triangular subdivision scheme for subdividing the geometric base in order to produce a refined triangular mesh, the wavelet coefficients being defined at the vertices of the triangles in the triangular mesh;

program code for defining a tree structure in which each node of the tree structure represents an associated triangle of the triangular mesh, there being a plurality of levels in the tree, the tree levels including a top level that includes at least one top level node, wherein each tree level below the top level includes a plurality of nodes that are children of parent nodes in a next higher level, wherein children nodes associated with a particular parent node represent triangles of the subdivided triangle associated with the particular parent node, wherein each wavelet coefficient at a particular level in the triangular mesh is uniquely defined at a particular vertex in the triangular mesh at that level and is uniquely assigned to a single one of the triangles at a next higher level of subdivision, such that each triangle at the next higher level of subdivision has in the range of zero to three assigned wavelet coefficients;

program code for setting a first threshold and outputting a stream of bits indicative of significant nodes, wherein a significant node is defined as a node that has at least one of, associated wavelet coefficient that exceeds the designated threshold and descendant node that has an associated wavelet coefficient that exceeds the designated threshold, wherein when a particular significant node is determined, the output stream of bits further includes bits indicative of the significance of children nodes of the particular significant node and the significant children are recursively examined for significant descendants and the output bit stream identifies the significant descendants.

19. A computer-implemented method for uncompressing data representing a function defined upon an M -dimensional manifold, the manifold being embedded in N -dimensional space where M and N are integers and N is greater than M , a geometric base being used as a coarse initial model of the manifold, the geometric base being subdivided into cells using a subdivision scheme to produce a refined mesh, the method comprising the steps of:

defining a tree structure in which each node of the tree structure represents an associated cell of the refined mesh, each node arranged to accommodate associated second generation wavelet coefficients, there being a plurality of levels in the tree, the tree levels including a top level that includes at least one top level node, wherein each tree level below the top level includes a plurality of nodes that are children of parent nodes in a next higher level, wherein children nodes associated with a particular parent node represent cells of a subdivided cell associated with the particular parent node;

setting a first threshold and inputting a stream of significance bits indicative of significant nodes, wherein a significant node is defined as a node that has at least one of, associated wavelet coefficient that exceeds the designated threshold and descendant node that has an associated wavelet coefficient that exceeds the designated threshold, wherein when a particular significant node is determined, the input stream of significance bits further includes bits indicative of the significance of children nodes of the particular significant node and the input bit stream identifies the significant descendants;

inputting a stream of data bits indicative of values for second generation wavelet coefficients, the wavelet coefficients being defined on cells within the refined mesh, wherein each wavelet coefficient at a particular level in the refined mesh is uniquely defined on a particular cell in the refined mesh at that level and is uniquely assigned to a single one of the cells, such that each cell may have assigned to it any number of wavelet coefficients; and

synthesizing the data representing the function from the wavelet coefficients in order to reconstruct the function defined upon the M-dimensional manifold.

20. A method as recited in claim 19 wherein each time a significant node is identified, an indicia of the sign of each significant coefficient associated with the significant node is also input.

21. A method as recited in claim 19 further comprising a refinement step that refines the setting and inputting steps, the refinement step including the substeps of:

setting a current threshold to a next threshold;

inputting a stream of bits indicative of newly significant nodes at the current threshold, wherein a newly significant node is defined as a node that is significant at the current threshold but was not significant at any previous threshold;

inputting a stream of data bits corresponding to the current threshold, the data bits indicative of values for the second generation wavelet coefficients.

22. A method as recited in claim 21 further comprising the step of repeating the refinement step wherein the current threshold is set to a threshold that is one-half of the previous threshold.

23. A method as recited in any of claims 19-22 wherein the M-dimensional manifold
5 embedded in N-dimensional space is a 2-dimensional manifold embedded in 3-dimensional space.

24. A method as recited in any of claims 19-23 wherein the subdivision scheme is a triangular subdivision scheme, the cells of the subdivided geometric base are triangles and the refined mesh is a refined triangular mesh that contemplates that each base triangle in a selected
10 level of the triangular mesh is divided into four subtriangles at the next lower level of subdivision, such that the vertices of the subtriangles include the vertices of the base triangle and the midpoints of the edges of the base triangle.

25. A method as recited in any of claims 19-24 wherein the synthesizing step uses a lifting scheme with a basis that is selected from the group consisting of linear, quadratic, cubic
15 and butterfly.

26. A method as recited in claim 24 wherein each wavelet coefficient is defined at the vertex of a triangle within the refined triangular mesh and the particular vertex at which each wavelet coefficient is defined is incident upon the midpoint of an edge belonging to two larger triangles at a next higher level of subdivision, and the wavelet coefficient is uniquely associated
20 with one of the two larger triangles.

27. A computer program product comprising a computer-usable medium having computer-readable code embodied thereon for uncompressing data representing a function defined upon an M-dimensional manifold, the manifold being embedded in N-dimensional space where M and N are integers and N is greater than M, a geometric base being used as a coarse
25 initial model of the manifold, the geometric base being subdivided into cells using a subdivision scheme to produce a refined mesh, the computer program product comprising the following computer-readable program code for effecting actions in the computer:

program code for defining a tree structure in which each node of the tree structure represents an associated cell of the refined mesh, each node arranged to accommodate associated
30 second generation wavelet coefficients, there being a plurality of levels in the tree, the tree levels including a top level that includes at least one top level node, wherein each tree level below the top level includes a plurality of nodes that are children of parent nodes in a next higher level, wherein children nodes associated with a particular parent node represent cells of a subdivided cell associated with the particular parent node;

program code for setting a first threshold and inputting a stream of significance bits indicative of significant nodes, wherein a significant node is defined as a node that has at least one of, associated wavelet coefficient that exceeds the designated threshold and descendant node that has an associated wavelet coefficient that exceeds the designated threshold, wherein when a particular significant node is determined, the input stream of significance bits further includes bits indicative of the significance of children nodes of the particular significant node and the input bit stream identifies the significant descendants;

program code for inputting a stream of data bits indicative of values for second generation wavelet coefficients, the wavelet coefficients being defined on cells within the refined mesh, wherein each wavelet coefficient at a particular level in the refined mesh is uniquely defined on a particular cell in the refined mesh at that level and is uniquely assigned to a single one of the cells, such that each cell may have assigned to it any number of wavelet coefficients; and

program code for synthesizing the data representing the function from the wavelet coefficients in order to reconstruct the function defined upon the M-dimensional manifold.

28. A computer-implemented method for uncompressing data representing a function defined upon a two-dimensional manifold, the manifold being embedded in three-dimensional space, a geometric base being used as a coarse initial model of the manifold, the geometric base being subdivided using a triangular subdivision scheme to produce a refined triangular mesh, the method comprising the steps of:

defining a tree structure in which each node of the tree structure represents an associated triangle of the triangular mesh, each node arranged to accommodate associated second generation wavelet coefficients, there being a plurality of levels in the tree, the tree levels including a top level that includes at least one top level node, wherein each tree level below the top level includes a plurality of nodes that are children of parent nodes in a next higher level, wherein children nodes associated with a particular parent node represent triangles of a subdivided triangle associated with the particular parent node;

setting a first threshold and inputting a stream of significance bits indicative of significant nodes, wherein a significant node is defined as a node that has at least one of, associated wavelet coefficient that exceeds the designated threshold and descendant node that has an associated wavelet coefficient that exceeds the designated threshold, wherein when a particular significant node is determined, the input stream of significance bits further includes bits indicative of the significance of children nodes of the particular significant node and the input bit stream identifies the significant descendants;

inputting a stream of data bits indicative of values for second generation wavelet coefficients, the wavelet coefficients being defined at the vertices of the triangles in the refined

triangular mesh, wherein each wavelet coefficient at a particular level in the triangular mesh is uniquely defined at a particular vertex in the triangular mesh at that level and is uniquely assigned to a single one of the triangles at a next higher level of subdivision, such that each triangle at the next higher level of subdivision has in the range of zero to three assigned wavelet coefficients; and

5 synthesizing the data representing the function from the wavelet coefficients in order to reconstruct the function defined upon the two-dimensional manifold.

29. A method as recited in claim 28 wherein each time a significant node is identified, an indicia of the sign of each significant coefficient associated with the significant node is also input.

10 30. A method as recited in claim 28 further comprising a refinement step that refines the setting and inputting steps, the refinement step including the substeps of:

setting a current threshold to a next threshold;

inputting a stream of bits indicative of newly significant nodes at the current threshold, wherein a newly significant node is defined as a node that is significant at the current threshold but was not significant at any previous threshold;

15 inputting a stream of data bits corresponding to the current threshold, the data bits indicative of values for the second generation wavelet coefficients.

31. A method as recited in claim 30 further comprising the step of repeating the refinement step wherein the current threshold is set to a threshold that is one-half of the previous threshold.

20 32. A method as recited in any of claims 28-31 wherein the manifold is a sphere and the geometric base is selected from the group consisting of an icosahedron, an octahedron and a tetrahedron.

33. A method as recited in any of claims 28-32 wherein the triangular subdivision scheme contemplates that each base triangle in a selected level of the triangular mesh is divided into four subtriangles at the next lower level of subdivision, such that the vertices of the subtriangles include the vertices of the base triangle and the midpoints of the edges of the base triangle.

34. A method as recited in any of claims 28-33 wherein the synthesizing step uses a lifting scheme with a basis that is selected from the group consisting of linear, quadratic, cubic and butterfly.

35. A method as recited in any of claims 28-34 wherein the particular vertex at which each wavelet coefficient is defined is incident upon the midpoint of an edge belonging to two larger triangles at a next higher level of subdivision, and the wavelet coefficient is uniquely associated with one of the two larger triangles.

5 36. A computer program product comprising a computer-usable medium having computer-readable code embodied thereon for uncompressing data representing a function defined upon a two-dimensional manifold, the manifold being embedded in three-dimensional space, a geometric base being used as a coarse initial model of the manifold, the geometric base being subdivided using a triangular subdivision scheme to produce a refined triangular mesh, the
10 computer program product comprising the following computer-readable program code for effecting actions in the computer:

program code for defining a tree structure in which each node of the tree structure represents an associated triangle of the triangular mesh, each node arranged to accommodate associated second generation wavelet coefficients, there being a plurality of levels in the tree, the
15 tree levels including a top level that includes at least one top level node, wherein each tree level below the top level includes a plurality of nodes that are children of parent nodes in a next higher level, wherein children nodes associated with a particular parent node represent triangles of a subdivided triangle associated with the particular parent node;

program code for setting a first threshold and inputting a stream of significance bits
20 indicative of significant nodes, wherein a significant node is defined as a node that has at least one of, associated wavelet coefficient that exceeds the designated threshold and descendant node that has an associated wavelet coefficient that exceeds the designated threshold, wherein when a particular significant node is determined, the input stream of significance bits further includes bits indicative of the significance of children nodes of the particular significant node and the input bit
25 stream identifies the significant descendants;

program code for inputting a stream of data bits indicative of values for second generation wavelet coefficients, the wavelet coefficients being defined at the vertices of the triangles in the refined triangular mesh, wherein each wavelet coefficient at a particular level in the triangular mesh is uniquely defined at a particular vertex in the triangular mesh at that level and is uniquely
30 assigned to a single one of the triangles at a next higher level of subdivision, such that each triangle at the next higher level of subdivision has in the range of zero to three assigned wavelet coefficients; and

program code for synthesizing the data representing the function from the wavelet coefficients in order to reconstruct the function defined upon the two-dimensional manifold.

1/15




Figure 1a (Prior Art)

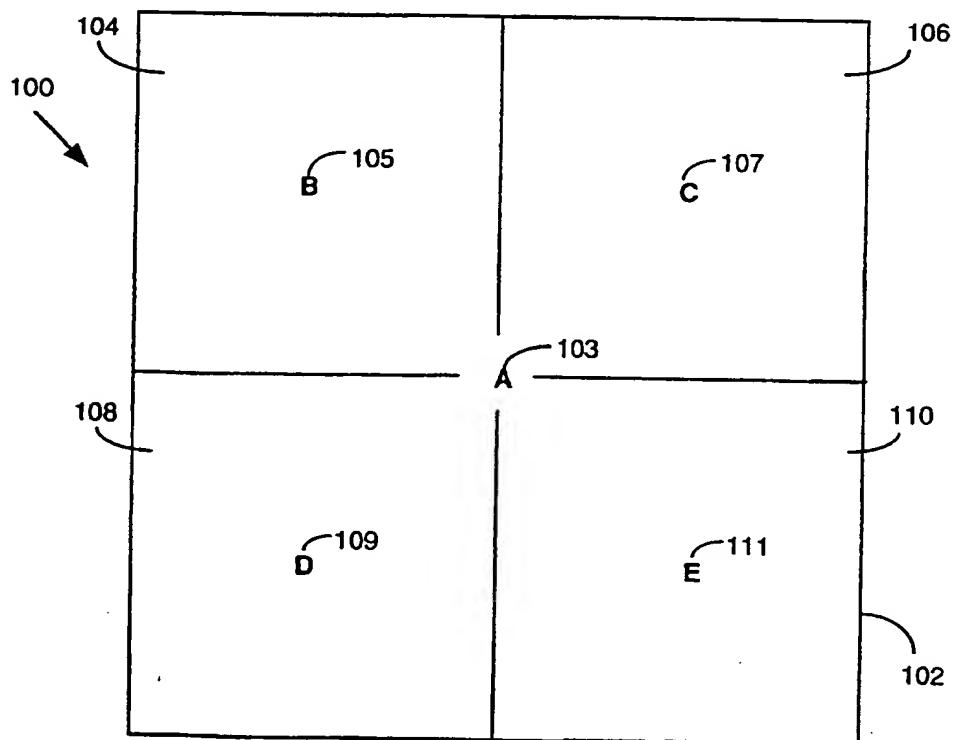


Figure 1b (Prior Art)

2/15




Figure 1c (Prior Art)




Figure 1d (Prior Art)

3/15




Figure 2a (Prior Art)

4/15

Figure 2b
(Prior Art)



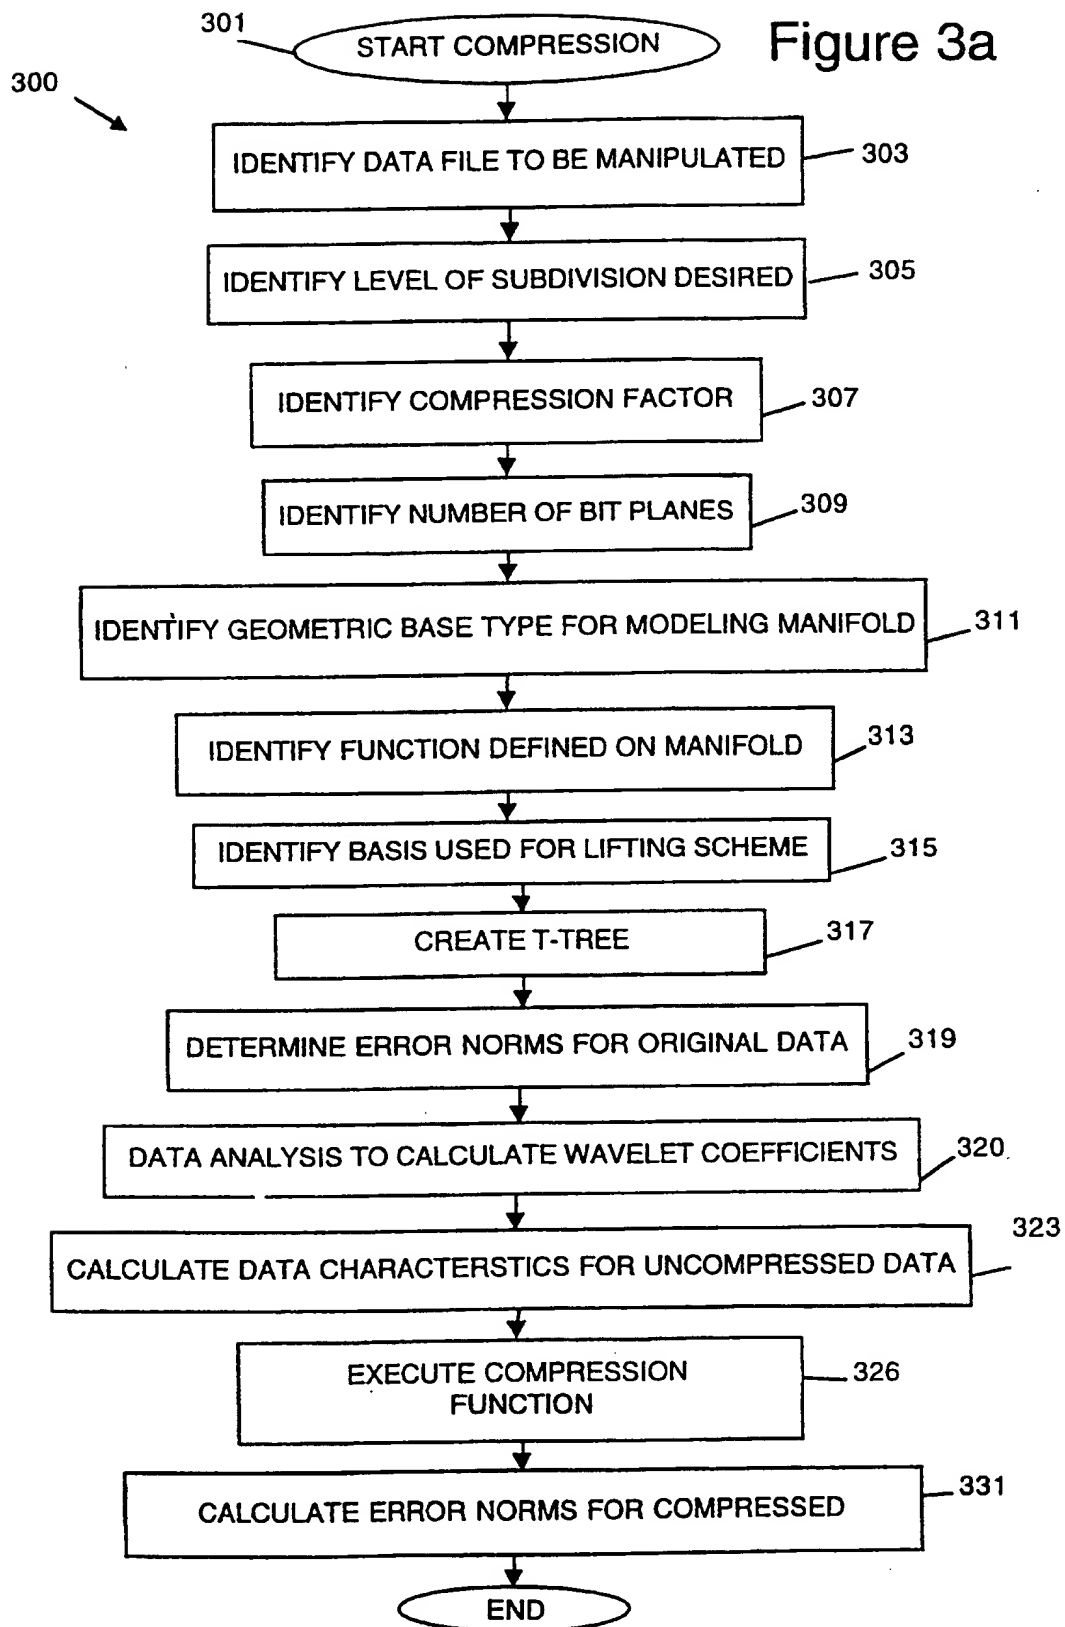


Figure 2c
(Prior Art)

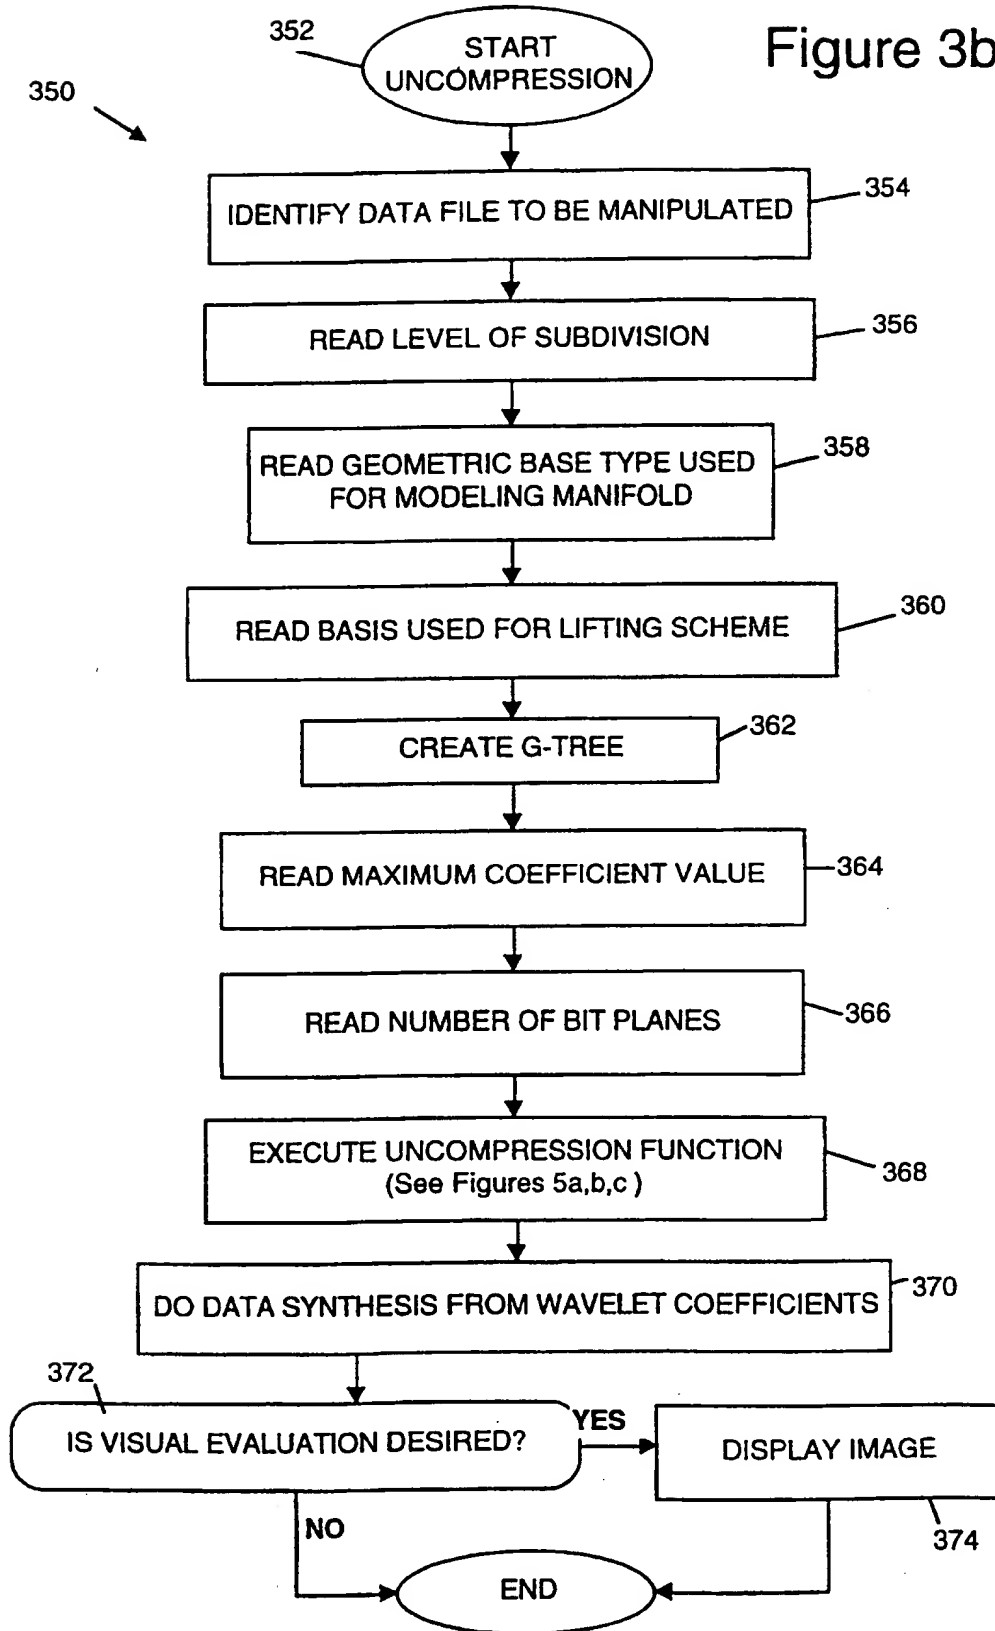


5/15



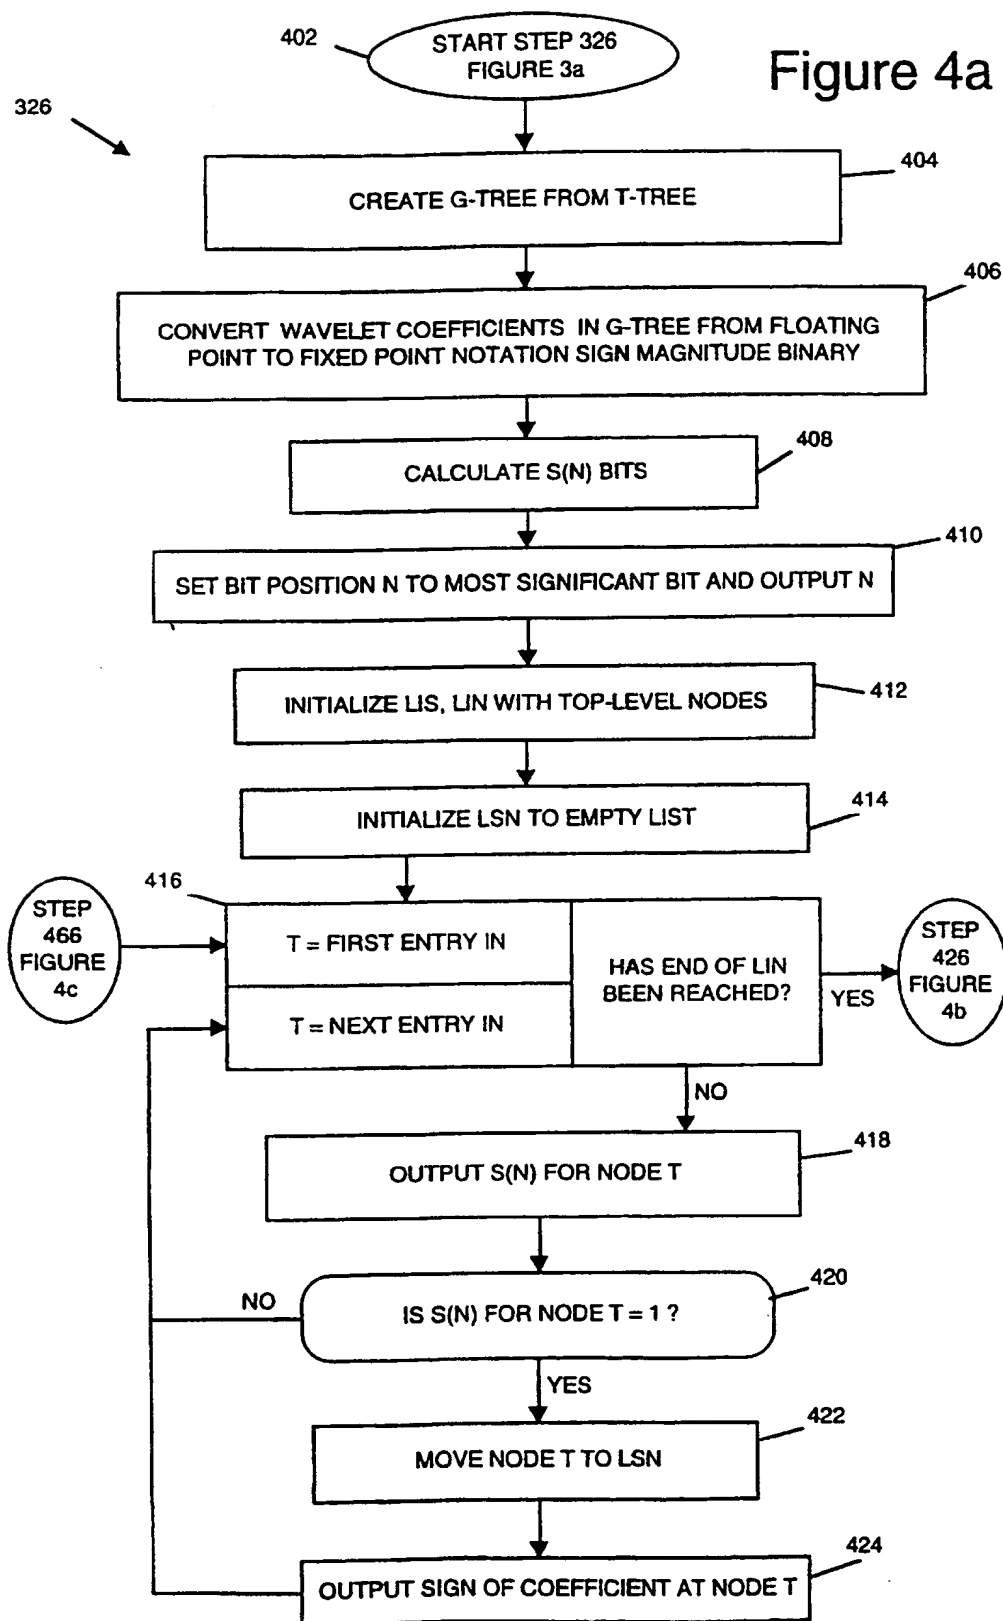
6/15

Figure 3b



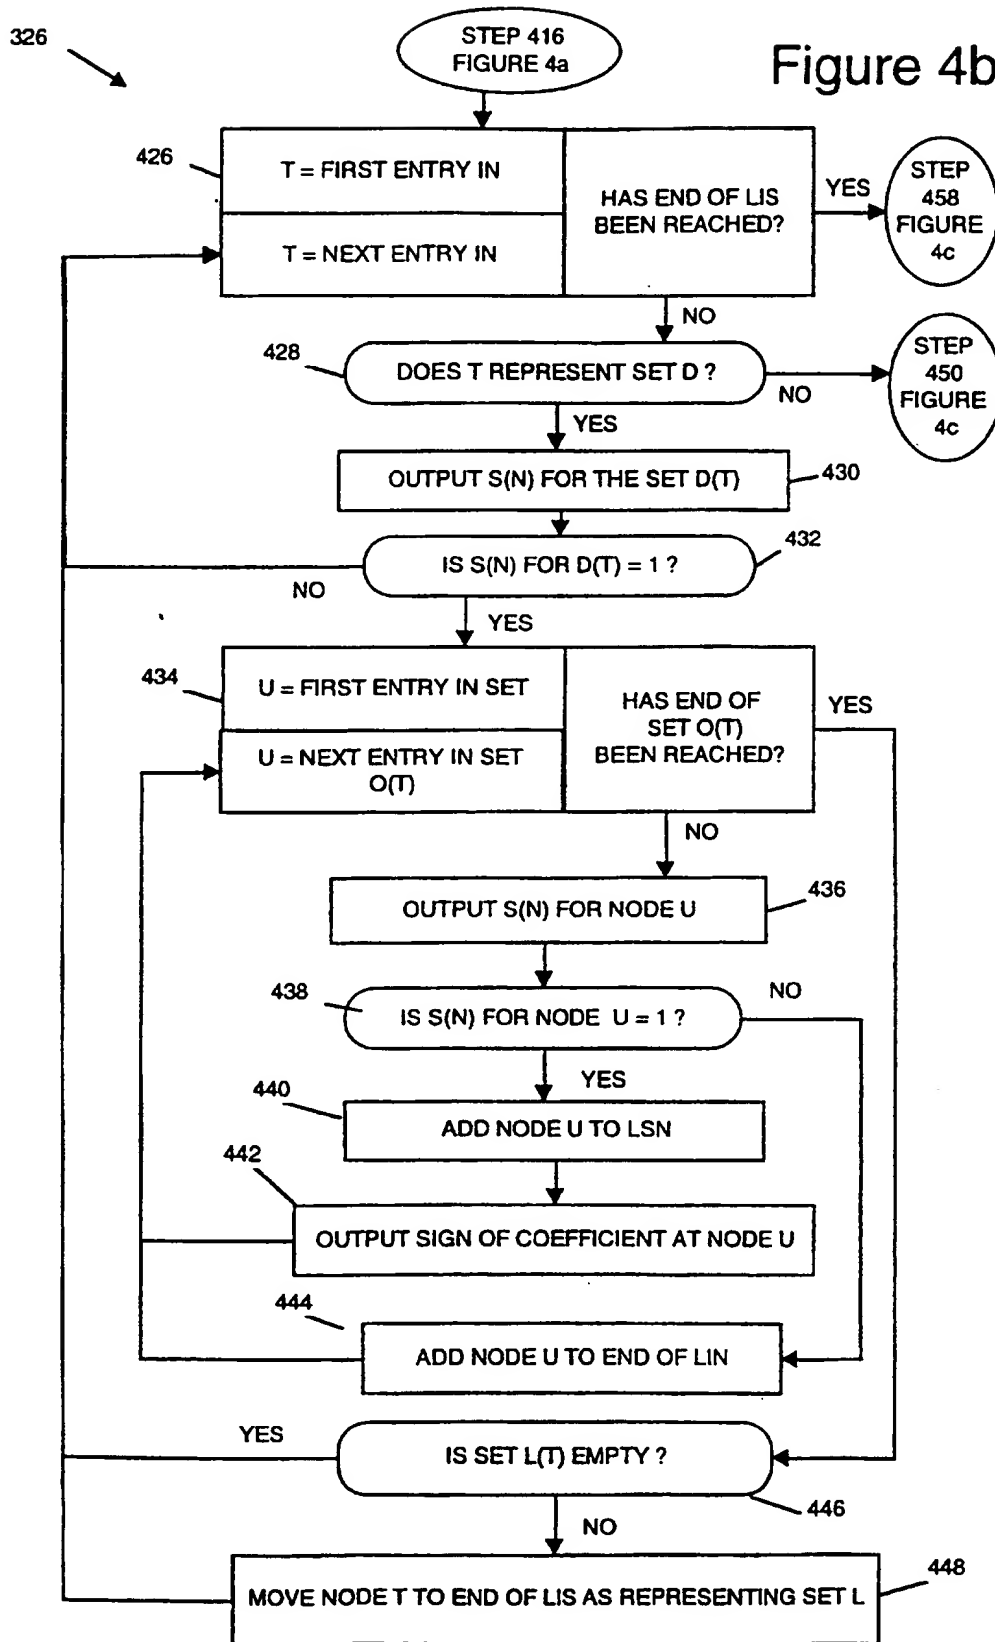
7/15

Figure 4a



8/15


Figure 4b



9/15

326

Figure 4c



10/15

Figure 5a

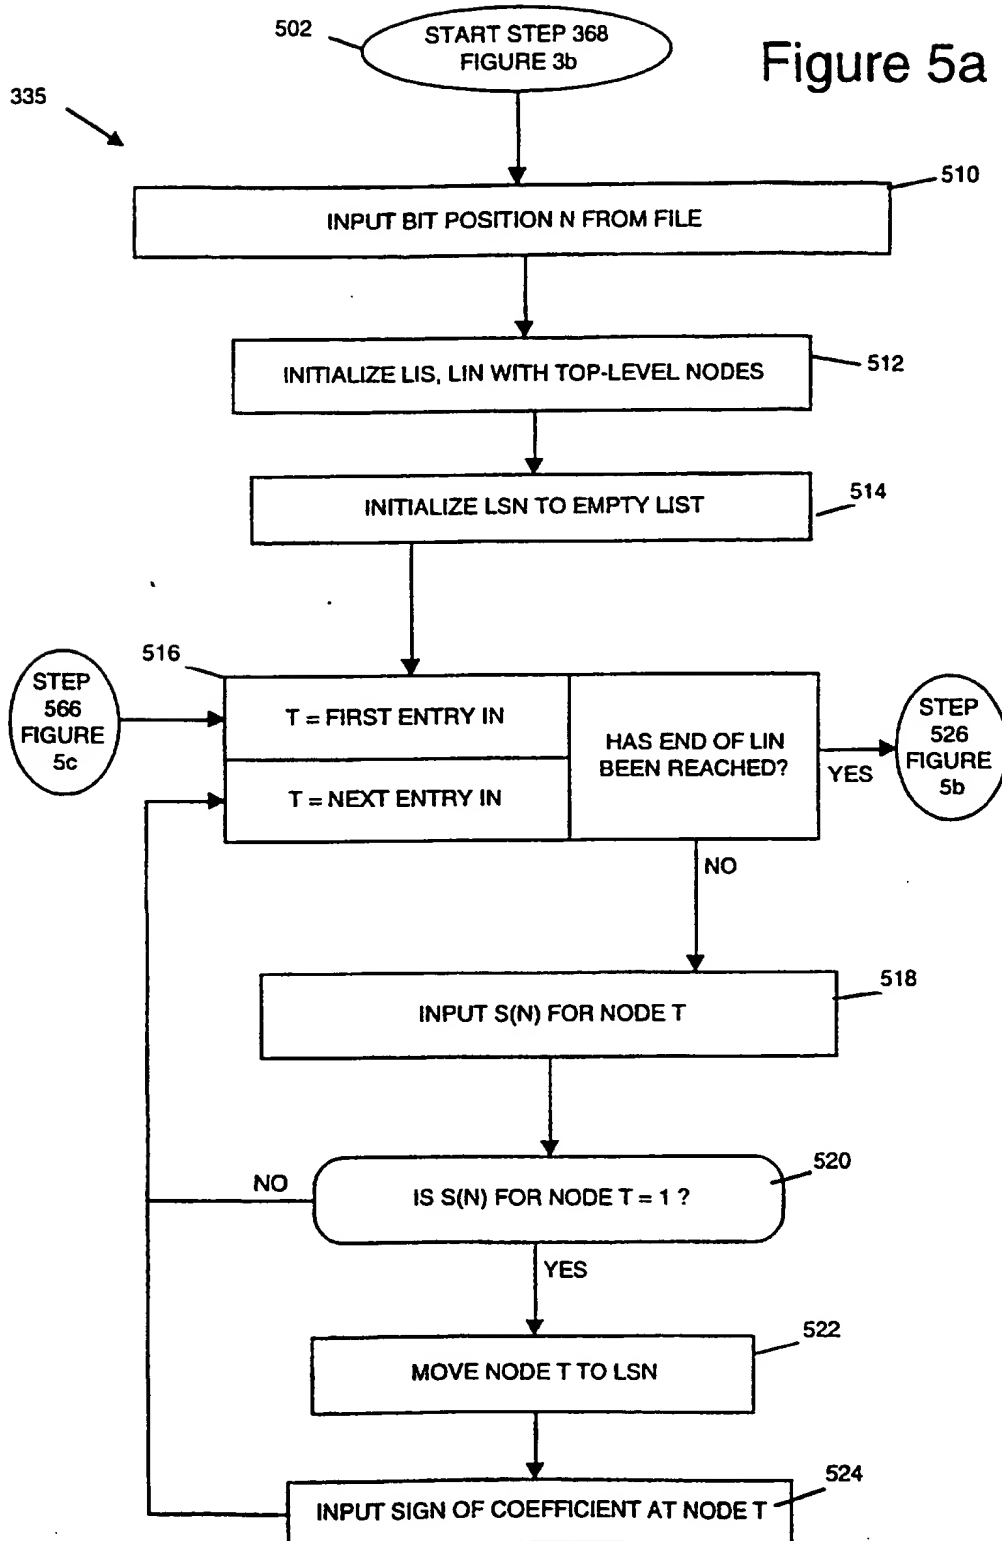
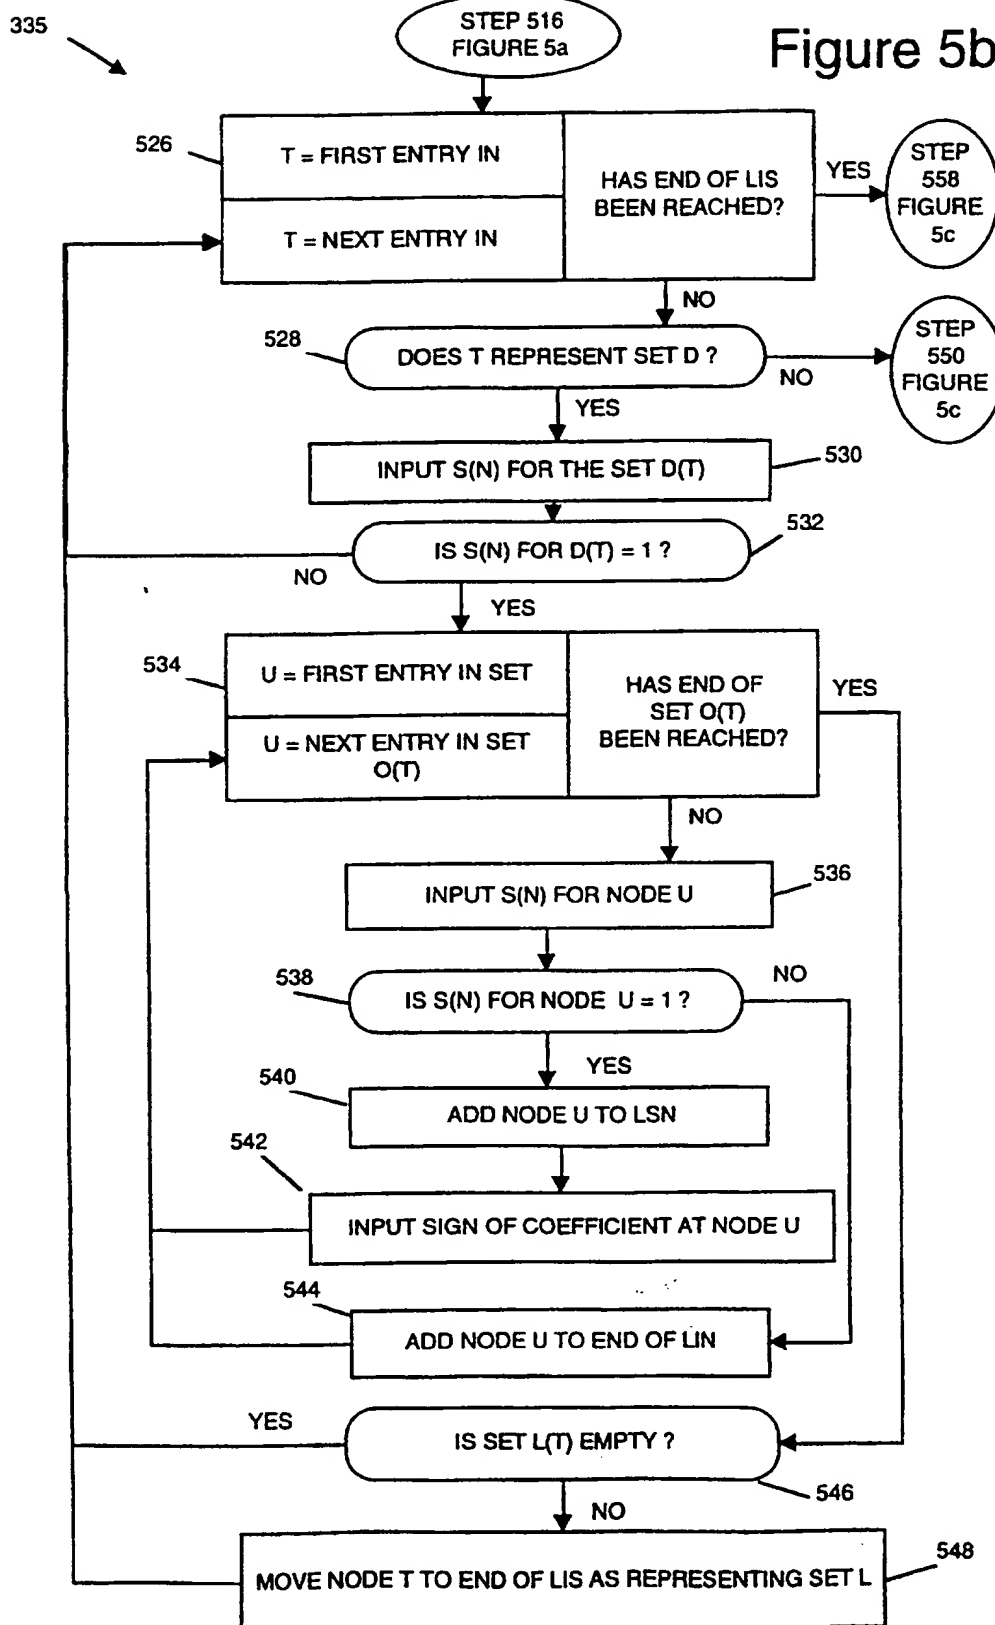



Figure 5b



12/15


335

Figure 5c



13/15

Figure 6



14/15




Figure 7a




Figure 7b




Figure 7c




Figure 8

INTERNATIONAL SEARCH REPORT

Int. Application No

PCT/US 97/03300

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 G06T9/00 G06T9/40

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 G06T

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>PROCEEDINGS OF THE VISUALIZATION CONFERENCE, ATLANTA, OCT. 29 - NOV. 3, 1995, 29 October 1995, INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, pages 135-142, XP000585401 GROSS M H ET AL: "FAST MULTIREOLUTION SURFACE MESHING" see the whole document ---</p> <p style="text-align: center;">-/--</p>	1-36

☒ Further documents are listed in the continuation of box C.

☐ Patent family members are listed in annex.

* Special categories of cited documents:

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

A document member of the same patent family

Date of the actual completion of the international search

19 June 1997

Date of mailing of the international search report

26.06.97

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Pierfederici, A

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US 97/03300

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	<p>PROCEEDINGS OF THE 4TH INTERNATIONAL SYMPOSIUM ON SPATIAL DATA HANDLING, ZURICH, SWITZERLAND, 23-27 JULY 1990, ISBN 3-906254-99-2, 1990, ZURICH, SWITZERLAND, UNIV. ZURICH, SWITZERLAND, pages 901-910 vol.3, XP000676005 DUTTON G: "Locational properties of quaternary triangular meshes" see the whole document</p> <p style="text-align: center;">---</p>	1-36
Y	<p>COMPUTER GRAPHICS PROCEEDINGS, LOS ANGELES, AUG. 6 - 11, 1995, 6 August 1995, COOK R, pages 161-172, XP000546225 SCHRODER P ET AL: "SPHERICAL WAVELETS: EFFICIENTLY REPRESENTING FUNCTIONS ON THE SPHERE" cited in the application see page 162, right-hand column, paragraph 6 - page 163, left-hand column, paragraph 2 see page 169, right-hand column, paragraph 7 - page 170, left-hand column, paragraph 4</p> <p style="text-align: center;">---</p>	1-36
A	<p>ACM TRANSACTIONS ON GRAPHICS, vol. 14, no. 4, 1 October 1995, pages 363-411, XP000558613 FLORIANI DE L ET AL: "HIERARCHICAL TRIANGULATION FOR MULTIREOLUTION SURFACE DESCRIPTION"</p> <p style="text-align: center;">---</p>	
A	<p>EXTRACTING MEANING FROM COMPLEX DATA: PROCESSING, DISPLAY, INTERACTION, SANTA CLARA, CA, USA, 14-16 FEB. 1990, vol. 1259, ISSN 0277-786X, PROCEEDINGS OF THE SPIE - THE INTERNATIONAL SOCIETY FOR OPTICAL ENGINEERING, 1990, USA, pages 242-253, XP000675767 FEKETE G ET AL: "Sphere quadrees: a new data structure to support the visualization of spherically distributed data"</p> <p style="text-align: center;">---</p>	
A	<p>ADVANCES IN SPATIAL DATABASES. THIRD INTERNATIONAL SYMPOSIUM, SSD '93 PROCEEDINGS, PROCEEDINGS OF SYMPOSIUM ON LARGE SPATIAL DATABASES (SSD'93), SINGAPORE, 23-25 JUNE 1993, ISBN 3-540-56869-7, 1993, BERLIN, GERMANY, SPRINGER-VERLAG, GERMANY, pages 510-529, XP000675774 OTOO E J ET AL: "Indexing on spherical surfaces using semi-quadcodes"</p> <p style="text-align: center;">-----</p>	